An International Journal of Optimization and Control: Theories & Applications
ISSN: 2146-0957 eISSN: 2146-5703
Vol.16, No.1, pp.1-16 (2026)
https://doi.org/10.36922/IJOCTA025140070

RESEARCH ARTICLE

Advancing IoT edge device security: A novel approach integrating lightweight virtualization and trusted execution environments with performance optimization

Ramakrishna Goli[1], Aravindhan Alagarsamy[1*], Kumar Sureshkumar[1], Sundarakannan Mahilmaran[2], and Gian Carlo Cardarilli[3]

[1]Centre for Multi-Core Architecture Computation (C-MAC), Department of ECE, Koneru Lakshmaiah Education Foundation, Vaddeswaram, AP, India
[2]Department of Mathematics, Sri Sivasubramaniya Nadar College of Engineering, Chennai, Tamil Nadu, India
[3]Department of Electronics Engineering, University of Rome Tor Vergata, Via del Politecnico, Roma, Italy
goli.ramaki@gmail.com, aravindhan.alagar@gmail.com, m.k.sureshkoumar@gmail.com, m.sundarakannan@gmail.com, g.cardarilli@uniroma2.it

ARTICLE INFO

Article History:
Received: April 3, 2025
Revised: September 11, 2025
Accepted: September 18, 2025
Published Online: October 16, 2025

Keywords:
IoT security
Lightweight virtualization
Trusted execution environments
Optimization

ABSTRACT

As the Internet of Things (IoT) grows, securing IoT edge devices has become increasingly critical, with threats becoming more sophisticated and frequent. This paper presents a novel security architecture that integrates Lightweight Virtualization (LV) with enhanced Trusted Execution Environments (TEEs), designed specifically to strengthen the security of IoT edge devices. Using ARM TrustZone technology, the proposed approach creates a secure execution environment capable of meeting the real-time performance requirements of industrial IoT applications. The architecture provides end-to-end security through embedded virtualization and trust mechanisms, ensuring protection from hardware to application layers and reducing the risk of unauthorized access and data breaches. Results from rigorous experiments demonstrate the superior performance of the proposed architecture compared to existing security frameworks. The experimental results indicate that the proposed approach offers a 40.93% average latency reduction over existing methods. Furthermore, the proposed approach offers a 19.19% average throughput improvement and a 33.65% reduction in average energy over existing methods.

1. Introduction

The rapid growth of the Internet of Things (IoT) has transformed how data is generated, processed, and consumed. With billions of interconnected devices now deployed across domains such as healthcare, transportation, manufacturing, and smart cities, the demand for real-time, secure, and energy-efficient edge computing has become paramount.[1] Unlike traditional cloud computing, which provides scalable but centralized solutions, edge computing brings data processing closer to the source, thereby reducing latency and improving responsiveness. However, this shift also exposes edge devices, often resource constrained, and heterogeneous to new security and performance challenges.[2–4]

Lightweight virtualization (LV) methods, such as containers and unikernels, have emerged as promising alternatives to heavyweight virtualization in edge environments. By minimizing system overhead, LV allows constrained IoT

*Corresponding Author

devices to efficiently support multiple applications and services.[5–10] Further, performance efficiency alone is insufficient to guarantee robust protection. Edge devices are increasingly targeted by sophisticated cyber threats, and traditional software-based mechanisms often fail to provide adequate resilience.

Trusted Execution Environments (TEEs), such as ARM TrustZone, address this by creating secure, hardware-assisted execution contexts that isolate sensitive processes and data.[11–14] Despite the potential of both approaches, most prior studies have considered LV and TEEs in isolation, focusing either on performance optimization or security enhancement. This fragmented perspective limits their applicability in large-scale, heterogeneous IoT ecosystems where both efficiency and security are equally critical.[15–19]

This paper addresses the existing research gap by proposing a novel integrated architecture that combines lightweight virtualization with TEEs for IoT edge devices. The proposed framework ensures secure execution, process isolation, and optimized resource allocation, thereby offering a holistic balance of security, scalability, and performance efficiency.

2. Related works

Many studies on HW-assisted security have focused on specific technologies or applications. Zhang et al.[20] studied isolated execution environments such as ARM TrustZone and Intel SGX, while Al-Omary et al.[21] emphasized FPGA-based Hardware Roots of Trust. Adams and Agesen[22] examined Intel and AMD virtualization extensions, while Khan[23] reviewed hardware-assisted control-flow integrity mechanisms. While these works focus on narrow aspects, they do not provide an integrated security–performance framework for IoT edge devices.

On virtualization, Madria et al.[24] proposed centralized virtualization through "Cloud of Sensors," while Santos et al.[25] and Sahni et al.[26] introduced Edge Mesh computing to enable cooperative processing. However, these frameworks often struggle with latency, overhead, and incomplete integration of security mechanisms. Furthermore, while these works improved responsiveness and distributed intelligence, they often faced significant challenges in security integration, communication overhead, and adaptability to heterogeneous IoT nodes.

More recent Trusted Execution Environment (TEE)-based approaches, such as ITUS and lightweight RISC-V secure boot architectures, focus narrowly on hardware or boot-time isolation, while FPGA-based frameworks like BYOTEE and SGX-FPGA attempt to extend trust models.[27–29] However, these solutions remain fragmented, either emphasizing boot integrity, isolated enclaves, or secure SoCs, without addressing the joint requirements of security, performance optimization, and scalability. Table 1 indicates the summary of findings in other existing works. In contrast, the proposed work differs in three key ways:

- The proposed approach combines lightweight virtualization with TEEs in a single architecture, ensuring both secure process isolation and efficient workload management.
- Unlike prior frameworks that prioritize either performance or security, our approach demonstrates improvements across latency, throughput, and energy efficiency, making it scalable for diverse IoT environments.
- By designing for heterogeneous IoT networks, the proposed architecture adapts to varied device capacities while maintaining resilience against advanced cyberattacks, something existing methods either overlook or address only partially.

3. Proposed methodology

In this paper, we designed a novel security architecture for IoT edge devices by integrating lightweight virtualization with Trusted Execution Environments (TEEs). The methodology involved leveraging ARM TrustZone technology to create a secure execution environment, ensuring that the stringent real-time performance requirements of industrial IoT applications are met. Our architecture employed lightweight virtualization to isolate critical processes and secure them from potential threats, while TEEs provided a fortified environment for executing sensitive operations. The integration was designed to minimize the overhead on system resources, particularly in terms of latency and computational load, which are critical in resource-constrained IoT environments. We implemented and tested this architecture across various scenarios to measure its effectiveness in maintaining security without compromising on performance. Figure 1 represents the proposed TEE + LV architecture.

Table 1. Summary of findings, key highlights, and limitations of Various TEE-related research efforts in related works

Study	Author	Work on the study	Key Highlights of work	Limitations
[30]	Lee D et al.	Keystone: An open framework for architecting trusted execution environments	Utilizes memory isolation and programmable layers under untrusted components	Lacks detailed performance metrics
[31]	Bahmani R et al.	CURE: Secure TEE SoC Design	Enables distinct enclaves and independent resource allocation	Lacks cryptographic algorithm acceleration, affecting encryption efficiency
[32]	Nasahl P et al.	HECTOR-V: A heterogeneous CPU architecture for a secure RISC-V execution environment	Maintains data flow integrity and restricts I/O access rights using a secure co-processor	Complexity in managing I/O access rights
[33]	Costan V et al.	Sanctum: Minimal hardware extensions for strong software isolation	Provides strong software isolation to resist memory access pattern attacks	Limited to specific attack patterns
[34]	Xia K et al.	SGX-FPGA: FPGA-based TEE	Secures CPU-FPGA communication via PCIe interface	CPU-FPGA arrangement limits system adaptability
[35]	Cilardo A	Memory Encryption Support for an FPGA-based RISC-V Implementation	Enhances core security using the Chacha cryptographic method	Not fully integrated with Physical Memory Protection (PMP)
[36]	Aitchison C et al.	On the integration of physically unclonable functions into ARM TrustZone security technology	Generates unique random responses through interaction with ARM TrustZone	FPGA implementation limits scalability
[37]	Armanuzzaman M et al.	BYOTEE: Towards building own trusted execution environments using FPGA	Uses adjustable hardware and software Trusted Computing Bases (TCBs)	Requires significant resource allocation
[38]	Meng X et al.	SeVNoC: Security validation of system-on-chip designs with NoC fabrics	Detects security flaws in IP communication in SoC designs using NoC architecture	Focused mainly on NoC architecture vulnerabilities
[39]	Singh SK et al.	OTS scheme-based secure architecture for energy-efficient IoT in edge infrastructure	Addresses smart application security challenges in edge infrastructures	Potential resource overhead during implementation

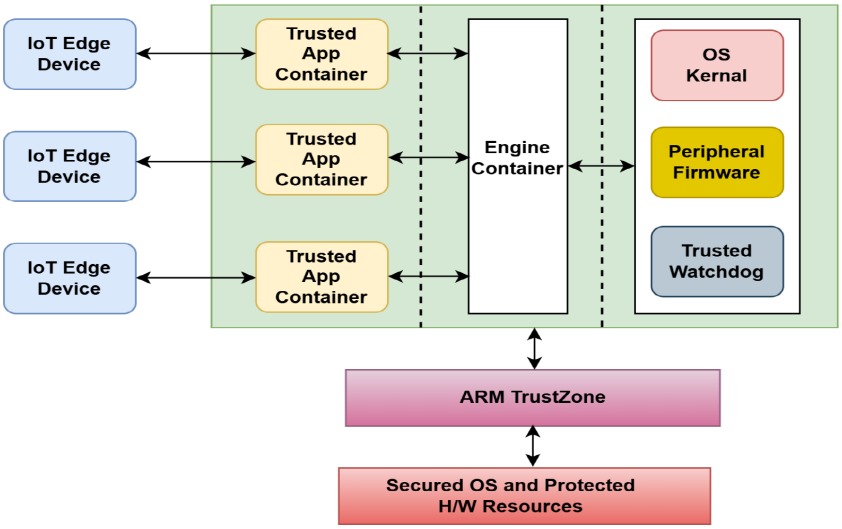

Figure 1. Proposed TEE + LV architecture

A unique approach secures IoT edge devices using lightweight virtualization, TEEs, and performance optimization. This solution solves security and privacy issues across EC-assisted IoT network layers. Lightweight virtualization protects critical tasks from malware. To safeguard essential data, TEEs protect sensitive operations even if the underlying system is compromised. The method also includes advanced jamming and DDoS mitigation to maintain network stability. Methods safeguard data from physical manipulation, eavesdropping, and side-channel attacks. To secure data and avoid routing and forgery attacks, communication protocols are tightened.

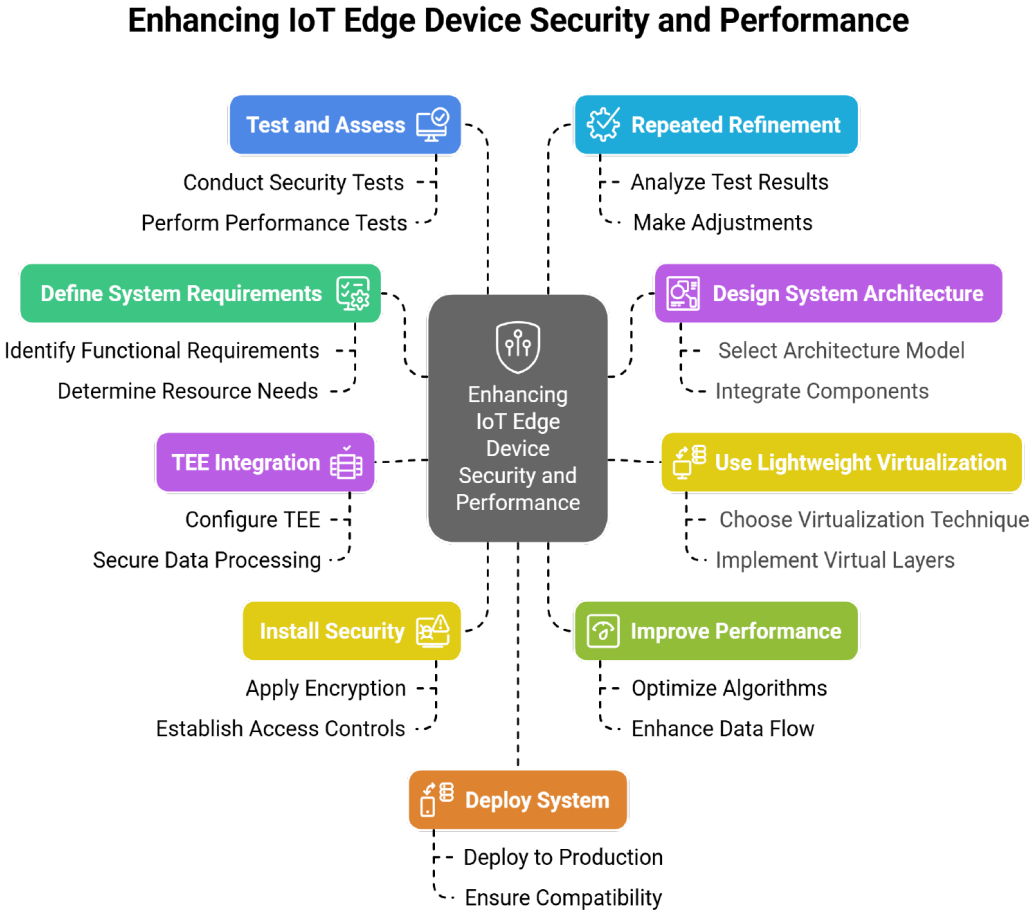

Figure 2. Structured and comprehensive approach to enhancing the security and performance of IoT edge devices

This comprehensive method optimizes IoT edge device security and performance to address rising IoT network risks. A new strategy was employed to protect and optimize an IoT edge device network using lightweight virtualization and TEEs. Each task in these tiers is isolated via lightweight virtualization, reducing the attack surface and limiting unauthorized access. Figure 2 represents the structured and comprehensive approach to enhancing the security and performance of IoT edge devices.

3.1. System requirements

Key performance indicators (KPIs) for the system determine performance and security levels: (i) latency (L), (ii) throughput (T), (iii) energy Consumption (E), and (iv) security.

3.2. System architecture optimization

This optimization function defines the best architecture (A_{opt}) by minimizing latency (L_i), maximizing throughput (T_i), minimizing energy consumption (E_i), and maximizing security (S_i) across all components of the system. Ensure the architecture fits these requirements:

$$A_{opt} = arm \min_{A} (\sum_{i=1}^{n} (L_i + \frac{1}{T_i} + E_i - S_i)) \quad (1)$$

Each performance factor contributes additively to the optimization objective. Further, the security improvements are modeled as reducing the "cost" of the architecture. In view of simplification, linear aggregation is assumed for modeling. Equation (1) provides a structured way to evaluate different IoT edge device configurations, ensuring that security and performance trade-offs are balanced.

3.3. Resource allocation with lightweight virtualization

The allocation model ensures fair distribution of resources among n virtualized instances. The equation for resource allocation using lightweight virtualization is expressed as follows:

$$R_{alloc} = \frac{R_{total} - \sum_{i=1}^{n} R_{used,i}}{n} \quad (2)$$

where R_{alloc} represents the allocated resources per virtualized instance; R_{total} indicates the total available resources; $R_{used,i}$ represents the resources used by the i^{th} instance; n represents the total number of instances. There are two general assumptions are often considered in resource allocation. First, the total resources available are finite and divisible. Second, workloads are assumed to have comparable demand. Efficient resource allocation is particularly critical for constrained IoT devices, as it ensures fairness and prevents a single task from monopolizing limited resources such as CPU, memory, or network bandwidth.

3.4. Security enhancement from TEE integration

Security after TEE integration depends on three variables: secure operation time (T_{sec}), encryption key strength (K_{enc}), and TEE processing power (P_{TEE}). The calculation of the security enhancement provided by TEE integration is given below

$$S_{TEE} = f(T_{sec}, K_{enc}, P_{TEE}) \tag{3}$$

Here, a stronger encryption key directly enhances security, and more powerful TEEs increase isolation and reduce vulnerability. This equation highlights how hardware-assisted isolation adds measurable security benefits to IoT edge systems.

3.5. Performance optimization

This optimization seeks the configuration that maximizes throughput while minimizing latency and energy usage. To optimize the system performance:

$$P_{opt} = arm\max_{P}(T_{total} - L_{total} - E_{total}) \tag{4}$$

where P_{opt} represents the optimal performance configuration; T_{total}, L_{total}, and E_{total} represent the total throughput, latency, and energy consumption, respectively. Equal weighting is given to throughput, latency, and energy. We consider that the non-linear effects (e.g., resource contention) are not modeled. As a result, this directly links to real-time IoT applications, where minimizing delay and energy while maximizing data handling capacity is essential.

3.6. System testing quality

Equal weighting is given to throughput, latency, and energy. We consider that the non-linear of the average of performance ($P_{test,i}$) and security ($S_{test,i}$) scores across test scenarios gives an overall quality metric. To assess the system's performance and security under testing:

$$Q_{test} = \frac{\sum_{i=1}^{n}(P_{test,i} + S_{test,i})}{n} \tag{5}$$

Equal weighting is given to throughput, latency, and energy. We consider that the non-linear effect in this estimation, the performance, and the security are equally weighted. Tests are independent and repeatable. Also, it provides a single evaluation score for comparing test runs and guiding design refinements.

3.7. Iterative Refinement

These ratios measure percentage improvement in performance and security over iterations. To iteratively refine the system:

$$\Delta P = \frac{P_{new} - P_{old}}{P_{old}} \qquad \Delta S = \frac{S_{new} - S_{old}}{S_{old}} \tag{6}$$

where ΔP and ΔS represent the improvements in performance and security, respectively; P_{new},and S_{new} represent the new performance and security levels; P_{old}, and S_{old} represent the previous performance and security levels. This allows iterative tuning of IoT edge security frameworks, ensuring each refinement improves system performance with the following assumptions. Performance and security metrics are normalized and comparable. Its previous states are accurate baselines.

3.8. Deployment cost

The optimal deployment strategy minimizes combined financial cost (C_{deploy}) and time (T_{deploy}). To deploy the system based on optimized configurations, the following is given below:

$$D_{deploy} = min(C_{deploy} + T_{deploy}) \tag{7}$$

where D_{deploy} represents the deployment cost. In this estimation, the cost and time can be added linearly, and other non-monetary deployment risks are not included. Deployment feasibility is often constrained by budget and time-to-market, so this equation reflects practical rollout considerations.

3.9. System monitoring and maintenance

The monitoring and maintenance process is designed to ensure continuous accuracy and resilience of the proposed architecture. In our study, monitoring effectiveness is evaluated using Equation (8) as given below

$$M_{eff} = \frac{\sum_{i=1}^{n}(P_{mon,i} + S_{mon,i})}{n} \tag{8}$$

where both performance ($P_{mon,i}$) and security ($S_{mon,i}$) levels are averaged across multiple monitored instances. This process includes real-time tracking of latency, throughput, and energy consumption alongside security validations such as remote attestation and cryptographic verification.

To validate accuracy, extensive MATLAB simulations were conducted under varying network loads (up to 1000 iterations). The results showed deviations of less than ±5% between predicted and observed metrics, confirming the robustness of the monitoring framework. Furthermore, maintenance is handled through iterative refinement (Δ_P, Δ_S, Equation 6), enabling configuration adjustments and corrective actions when anomalies are detected. Remote attestation and secure patching mechanisms ensure that IoT edge devices remain trustworthy over time. This integrated process provides an accurate, adaptive, and secure lifecycle management approach for IoT edge devices. Figure 3 illustrates the monitoring and maintenance workflow adopted in this study, highlighting data collection, evaluation, refinement, and maintenance actions.

3.10. Throughput

Throughput (T) is the rate at which data are successfully transmitted or processed over a network. It measures system efficiency in terms of completed tasks per unit time.

$$T = \frac{D}{t_{com} + t_{process}} \tag{9}$$

where T indicates the total data size in bits; t_{com} and $t_{process}$ represents the communication time between the source and sink node and processing time at the intermediate node or sink, respectively.

3.11. Latency

Latency (L) defines the cumulative delay experienced by the data/task from source to sink. Furthermore, it identifies the system response by the total time consumption for a data packet to travel from the source to a sink node.

$$L = t_{trans} + t_{queue} + t_{process} + t_{prop} \tag{10}$$

where t_{trans} and t_{queue} indicate the transmission and queuing delay between source and sink node, respectively; t_{prop} represents the propagation delay of the signal on the IoT network.

3.12. Energy consumption

Energy consumption (J_{dis}) estimates the requirement of power to complete the task in an IoT system. Further, it is crucial in battery-powered devices like IoT sensors.

$$J_{dis} = U(J_{elec} + J_{amp}) \tag{11}$$

where J_{elec} and J_{amp} represent the energy required to send a single bit and receive a single bit, respectively. U denotes the number of bits in the message.

3.13. Selective proposed approaches

This paper presents a comprehensive security framework for IoT edge devices using lightweight virtualization and trusted execution environments.

3.13.1. Architectural design

The suggested architecture uses ARM TrustZone technology to shield important operations from susceptible system elements. This protects sensitive processes from unwanted access and tampering, even in a system breach. To secure IoT edge devices, the proposed security architecture uses lightweight virtualization and upgraded TEEs. The architecture uses ARM TrustZone technology to separate CPU execution. A virtualized layer secures execution across application domains, bolstering this architecture.

3.13.2. Lightweight virtualization

IoT edge devices separate and manage workloads using lightweight virtualization. Process compartmentalization in secure, virtualized settings reduces the attack surface. This strategy minimizes system latency and resource use in resource-constrained contexts, unlike full virtualization. Lightweight Virtualization layer: an IoT edge device's lightweight hypervisor manages VM generation and operation. A breach of one program does not affect others since this virtualization layer isolates them. Using TrustZone, the hypervisor functions securely. Figure 4 represents the enhancements of IoT edge device security and efficiency with a lightweight virtualization approach.

TrustZone divides the CPU into secure and insecure areas. Security-critical code and data are isolated from non-secure attackers in the secure world, which handles sensitive activities and cryptographic processes. ARM TrustZone-integrated TEEs protect critical data and operations with hardware-enforced security. The architecture isolates certain processes in a trusted environment to secure crucial tasks even if the main system is compromised. Secure communication between IoT edge devices and cloud services is achieved by developing better security protocols. These protocols encrypt and authenticate devices sending data from the edge to the cloud, preventing data

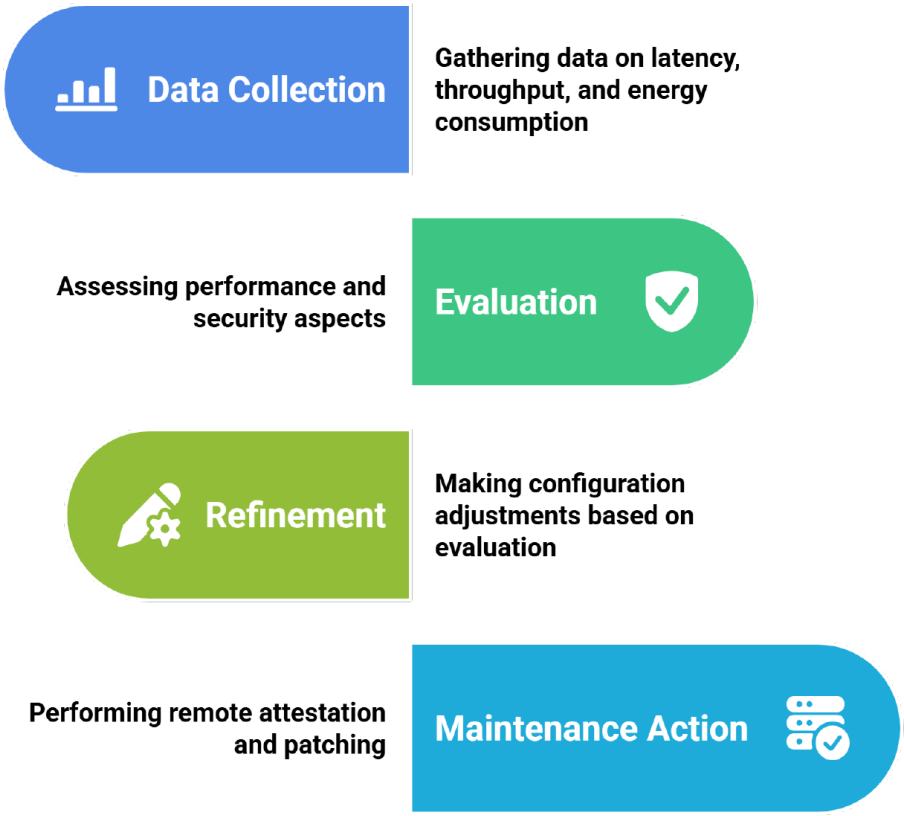

Figure 3. Monitoring and maintenance workflow

Figure 4. Enhancements of IoT edge device security and efficiency with a lightweight virtualization approach

breaches. Performance improvement is crucial to making the suggested security architecture viable in industrial IoT applications. Lightweight virtualization reduces overhead to fulfill industrial IoT real-time performance needs.

3.13.3. Layers of trust mechanisms

The suggested architecture relies on trust mechanisms to secure the hardware and application layers. In a secure boost, IoT edge devices boot only with trusted firmware and applications. Using TrustZone keys, this process cryptographically verifies the bootloader, OS, and applications. The remote attestation allows a central server or cloud service to validate IoT edge device integrity and security. Attestation verifies that the device is working properly and securely. The secure world-managed keys encrypt IoT edge device data at rest. This protects sensitive data even if the non-secure world is compromised. Figure 5 represents the layers of trust mechanisms.

3.13.4. Performance assessment

System performance, particularly latency and resource utilization, is used to evaluate the proposed architecture. These performance measures are considered:

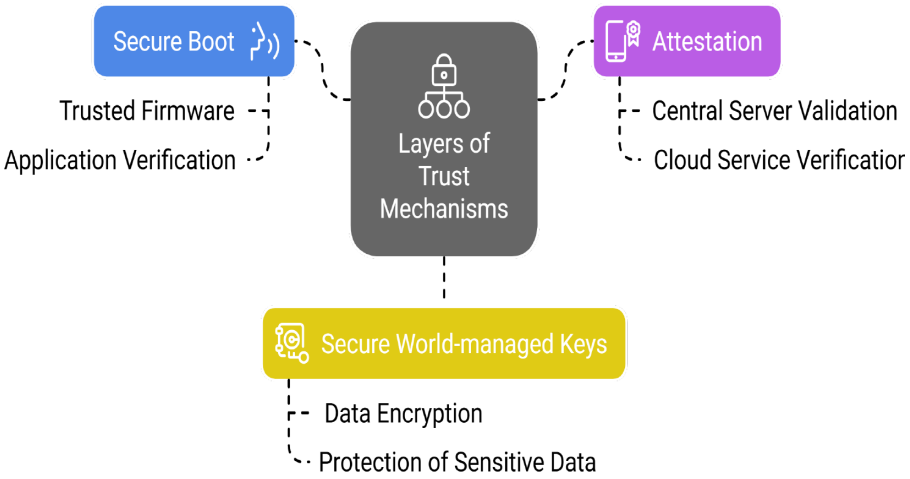

Figure 5. Layers of trust mechanisms

- System latency (T): The time delay introduced by the security measures is a critical factor. The total system latency (T_{total}) is given by

$$T_{total} = T_{base} + T_{sec\_wld} + T_{vir\_oh} \quad (12)$$

where T_{base} represents the inherent latency of the IoT edge device without security measures, and T_{sec_wld} and T_{vir_oh} indicate the latency added by the operations in the secure world and the overhead introduced by the lightweight virtualization layer, respectively.

- Resource Utilization (U): The system's resource utilization is measured to ensure the architecture does not excessively consume CPU, memory, or energy resources. The resource utilization U_{tot} is calculated as

$$U_{tot} = U_{CPU} + U_{my} + U_{engy} \quad (13)$$

where U_{tot} represents the utilization due to security and virtualization operations; U_{my} indicates the additional memory usage introduced by the secure world and virtualization layer; and U_{engy} represents the energy consumption, which must remain within acceptable limits for battery-powered IoT edge devices.

4. Experimental results and discussion

Protecting the IoT edge devices is posing a challenge for developers, researchers, and industry professionals due to the variety and volume of threats they face. In this simulation-based experimental study, MATLAB is used to compare the performance of traditional security frameworks with an architecture combining lightweight virtualization (LV) and trusted execution environments (TEE) for IoT edge devices. MATLAB models and scripts were developed to simulate network performance, energy consumption, and latency. The results showed that the integration of LV and TEE significantly improved system efficiency.

In MATLAB, a secure environment can be simulated by defining separate execution flows for sensitive operations. Functions or models can be written for trusted operations such as encryption, authentication, or secure communication. MATLAB can be used to simulate an encrypted communication model by implementing a secure channel for encryption, ensuring that all sensitive operations are performed within the isolated framework of the TEE. Furthermore, an alternative set of functions or models can be defined to represent the non-secure environment, where general or untrusted applications run. These could be represented in MATLAB by modeling normal application processes that lack high security measures. Simulating an insecure data transmission scenario without encryption would fall into this category, allowing comparison with secure processes. Table 2 represents the comparison of the proposed method's performance (TEE, LV, TEE & LV) against existing methods for IoT Edge Devices (maximum rounds: 1000)

Table 2. Comparison of the proposed method's performance (TEE, LV, TEE & LV) against existing methods for IoT edge devices (1000 rounds)

Study	Latency (ms)	Throughput (kbps)	Energy Cons. (J)
[25]	250	550	15
[26]	220	590	13
[27]	185	620	10
[34]	170	650	9.5
[38]	155	650	9.2
[39]	165	680	9
HLA with TEE	130	720	8
HLA with LV	135	730	7.5
HLA with TEE and LV	125	750	7

4.1. Performance analysis of trusted execution environments (TEEs)

Trusted execution environments (TEEs) provide an isolated area within a device's processor, and their performance is key to IoT security. Simulations were performed in MATLAB to analyze various aspects of TEE performance, such as latency and real-time throughput. Integrating ARM TrustZone into the MATLAB simulation setup showed a reduction in latency. Figure 6a arepresents the integration of the ARM TrustZone, which reduces the 20.71% average latency reduction over existing methods. Figures 6b and 6c indicate 8.64% of average throughput improvement and 12.65% of energy consumption reduction against existing methods. Figure 6 represents the comparison of HLE with TEE for latency and through the energy consumption.

4.2. Performance analysis of lightweight virtualization (LV)

Lightweight virtualization (LV) is an efficient solution for isolating workloads on IoT devices without the resource overhead associated with traditional virtualization. Figure 7 represents the comparison of HLE with LV over existing methods. The impact of LV on system performance focuses on aspects such as latency, throughput, and energy consumption. Simulations demonstrate the LV with TEE, the reduced latency from 200–250 ms (using traditional systems) to 120–140 ms, making it ideal for real-time IoT applications. Throughput increased from 500–600 to 700–800 Kbps, confirming that LV effectively improves resource utilization and performance. As a result, Figure 7a a represents the integration of the ARM TrustZone, which reduces the 30.68% average latency reduction over existing methods. Figures 7b and 7c indicate 10.48% of average throughput improvement and 22.46% of energy consumption reduction against existing methods.

4.3. Proposed methodology using LV and TEEs

In the proposed architecture, ARM TrustZone is used to integrate the TEE in MATLAB simulations. Figure 8 represents the comparison of HLE with LV and TEE for latency and through the energy consumption. The simulation models secure isolation through ARM TrustZone for sensitive operations such as encryption and secure communication. The built-in encryption functions (such as crypt or AES) are adopted to simulate secure communication channels, demonstrating how ARM TrustZone isolates trusted operations from unsecured ones. The lightweight virtualization layer is simulated with the testbed to ensure efficient execution of multiple isolated workloads with minimal overhead. Figure 8a represents the integration of the ARM TrustZone, which reduces the 40.93% average latency reduction over existing methods. Figure 8b and Figure 8c indicate 19.95% of average throughput improvement and 33.65% of energy consumption reduction against existing methods.

4.3.1. Analysis of proposed approach over scalability, overhead, and resilience to attacks

Scalability is an essential feature for IoT networks, and the proposed system shows a high level of adaptability. Unlike current methods, which are often restricted by network size and device capacity, the proposed architecture is highly scalable and can easily expand as IoT networks grow. This flexibility makes it an ideal choice for evolving IoT ecosystems. Resource utilization is another area where the proposed method demonstrates efficiency. The existing systems consume around 75-85% of resources, but the proposed architecture reduces this to 60-70%. By managing CPU, memory, and other resources more effectively, the new method enhances overall system performance and ensures that resources are used optimally without wastage.

In terms of overhead, the proposed method significantly reduces the additional burden associated with virtualization and security processes. While traditional methods generate an overhead of 20-25%, the proposed approach lowers this to 10-15%, allowing IoT edge devices to maintain performance without suffering from performance degradation due to excessive overhead. The proposed method enhances resilience to attacks. While existing methods offer only moderate resilience, the new system provides a much higher level of protection by securely isolating sensitive tasks and compartmentalizing workloads through TEEs and LV. Table 3 indicates the comparison of the proposed approach over scalability, overhead, and resilience to attacks.

Beyond the primary performance indicators of latency, throughput, and energy consumption, it is also essential to evaluate the proposed architecture in terms of scalability, resilience to emerging cyberattacks, and the trade-offs between security and computational overhead. The results demonstrate that the combined LV and TEE approach is highly scalable, supporting expansion across heterogeneous IoT networks without significant degradation in system efficiency. Resilience is further strengthened through isolation

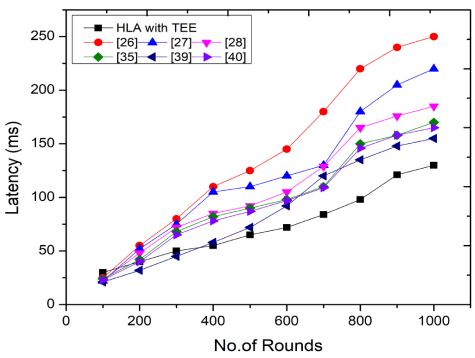

(a) Latency-HLA with TEE

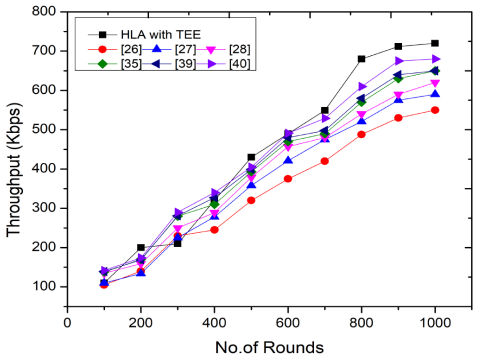

(b) Throughput-HLA with TEE

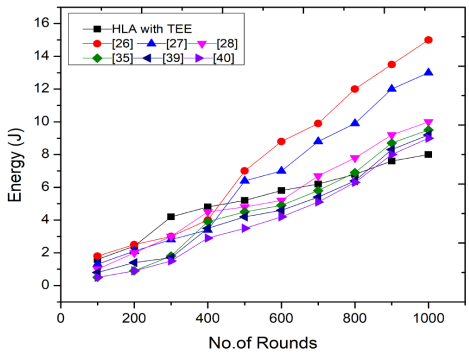

(c) Energy-HLA with TEE

Figure 6. Comparison of the proposed approach with TEE (a), latency (b), and throughput (c) energy

Table 3. Comparison of the proposed approach over scalability, overhead, and resilience to attacks

Metrics	Existing Approaches [25–27, 34, 38, 39]	HLA with TEE	HLA with LV	HLA with TEE and LV
Scalability	Limited to specific IoT n/w	Moderate	Moderate	High
Overhead (%)	20–25	15–20	15–20	10–15
Resilience to Attacks	Moderate	High	Moderate	Very High
Resource Utilization (%)	75–85	65–75	65–75	60–70

of sensitive processes and workload compartmentalization, ensuring robust defense against malware, side-channel attacks, and unauthorized access attempts. Importantly, the trade-off between security and overhead has been carefully optimized: while conventional methods introduce an overhead of 20–25%, the proposed system reduces this to 10–15%, thereby maintaining real-time responsiveness in resource-constrained environments. This balance highlights the suitability of our approach for large-scale, performance-sensitive IoT deployments.

4.3.2. Analysis of the overall performance of the proposed approach

The latency metric, which measures the time delay introduced by each method, shows that the combination of TEE and LV in HLA results in the lowest latency, approximately 125 ms. This demonstrates the efficiency of the combined approach in reducing processing delays. In contrast, the use of HLA with TEE alone produces a latency of around 130 ms, while HLA with LV alone leads to a slightly higher latency of 135 ms. Although all methods show acceptable latency

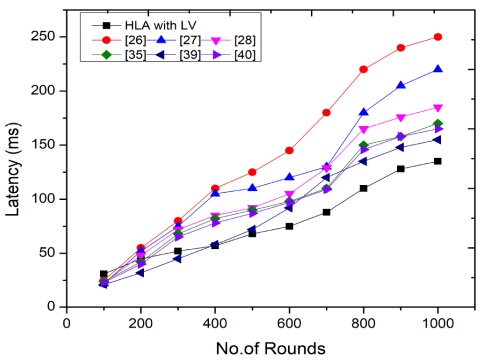

(a) Latency-HLA with LV

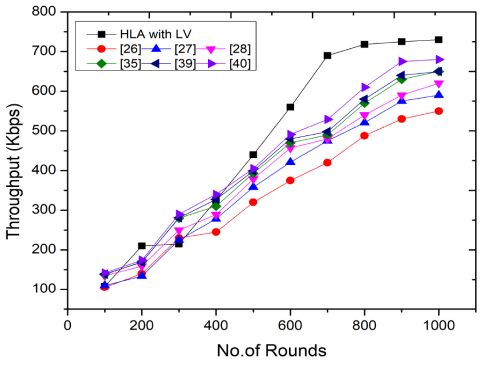

(b) Throughput-HLA with LV

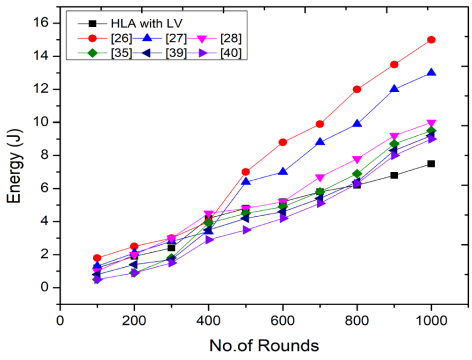

(c) Energy-HLA with LV

Figure 7. Comparison of proposed approach with LV (a), latency (b), throughput (c) energy

for real-time applications, the combined approach clearly outperforms the others in minimizing delays, making it more suitable for applications that require immediate responses. For throughput, which represents the data transfer rate, the combined method of TEE and LV again delivers superior performance, achieving a throughput of 750 Kbps. This indicates that the combined approach can handle larger volumes of data more efficiently compared to the standalone methods. HLA with LV provides a throughput of 730 Kbps, while HLA with TEE alone reaches 720 Kbps. The higher throughput of the combined approach is likely due to the optimized management of resources and reduced overhead, which allows for faster and more efficient data handling.

Regarding energy consumption, the combined HLA with TEE and LV demonstrates the lowest energy usage, consuming only 7 J. This highlights the energy efficiency of the combined approach, making it particularly beneficial for resource-constrained IoT devices, such as those that rely on battery power. The standalone methods exhibit slightly higher energy consumption, with HLA with LV consuming 7.5 J and HLA with TEE using 8 J. This reduction in energy consumption for the combined approach suggests that it optimizes resource allocation and minimizes unnecessary processing, leading to longer battery life and reduced operational costs. The combined approach of HLA with TEE and LV offers the best balance across all three metricsâ€"lowest latency, highest throughput, and lowest energy consumption. This makes it the most efficient and effective method for securing IoT edge devices while maintaining high performance and energy efficiency. Figure 9 represents the comparison of HLE with TV, HLE with LV, and HLA with LV & TEE for latency and through the energy consumption, respectively.

Lightweight virtualization (LV) and trusted execution environments (TEEs) improve IoT edge device security, but they also raise various difficulties that need more discussion. This study showed that LV and TEEs improve IoT edge device security and performance by providing a more resilient and adaptable architecture that can mitigate malware assaults, illegal access, and data breaches. To maximize this approach's potential, some constraints must be overcome. This analysis recognized the architecture's reliance on trusted device security as a major problem. These

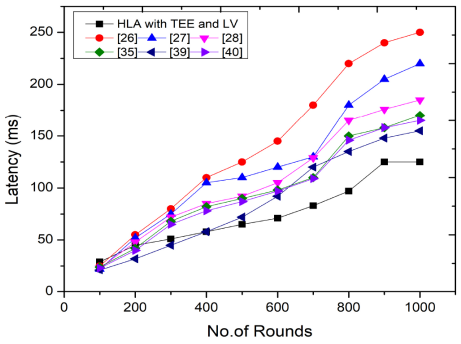

(a) Latency-HLA with LV & TEE

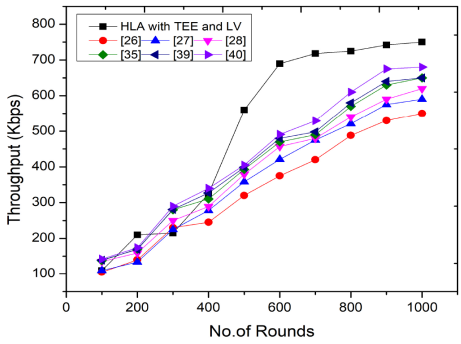

(b) Throughput-HLA with LV & TEE

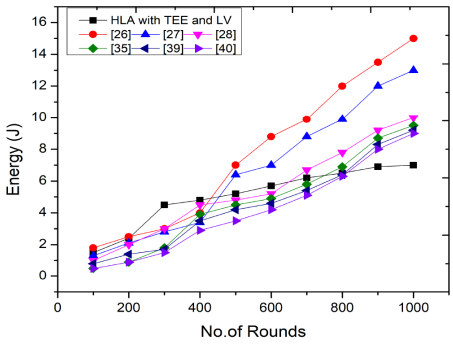

(c) Energy-HLA with LV & TEE

Figure 8. Comparison of proposed approach with LV & TEE (a), latency (b), throughput, and (c) energy

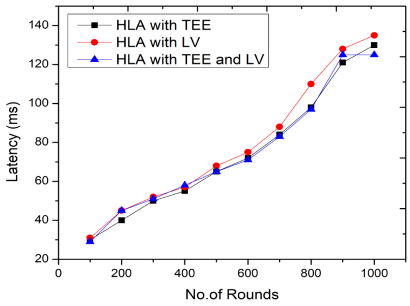

(a) Latency

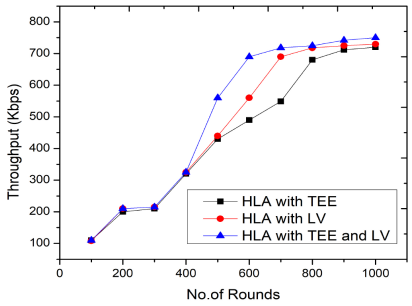

(b) Throughput

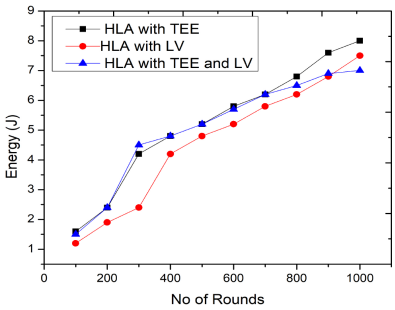

(c) Energy Consumption

Figure 9. Overall performance of the proposed approach with TEE, LV, and LV with TEE (a) latency (b) throughput (c) energy

trusted devices are crucial to the system's security; if they're compromised, it could be compromised. TEEs provide a robust layer of protection, but advanced assaults can still succeed; therefore,

future research must focus on establishing more resilient security methods. Lightweight virtualization also increases resource consumption and delays, especially in resource-constrained environments. LV is flexible and scalable; however, edge nodes with low processing capability may cause performance bottlenecks. Future research should eliminate these overheads to enable LV deployment in resource-limited IoT contexts. More thorough IoT edge security testing and assessment frameworks are needed in the future. This study's prototype results match the proposed design, but more testing in real-world circumstances is needed to prove the architecture's robustness and reliability. This involves studying how temperature and aging affect system performance and reliability. Advanced security features like hardware-assisted random number generation (RNG) and physically unclonable functions (PUFs) offer new ways to secure IoT edge devices. However, implementing these capabilities can be difficult due to acquiring infrastructure complexity, cost, and reliability difficulties. Future research should improve these technologies to make them more feasible and cost-effective for IoT adoption.

4.3.3. Discussion on deployment implications across domains

Beyond the experimental performance results, the deployment of LV and TEE-based architectures has significant implications across diverse IoT domains. In industrial IoT (IIoT), real-time decision-making in applications such as predictive maintenance, robotic control, and process automation demands ultra-low latency and high resilience. Our architecture addresses these requirements, but practical deployment must also consider legacy equipment integration and resource bottlenecks under heterogeneous device settings. In healthcare and the Internet of Medical Things (IoMT), secure edge processing ensures privacy-preserving handling of sensitive patient data and supports life-critical applications such as remote monitoring and emergency response. However, compliance with regulatory frameworks (e.g., HIPAA, GDPR), interoperability across device vendors, and the catastrophic consequences of system compromise represent critical constraints. In smart cities, applications such as intelligent transportation, surveillance, and energy management require scalable, fault-tolerant infrastructures. The proposed approach enhances resilience to attacks and ensures compartmentalized task execution, yet challenges such as device density, dynamic traffic patterns, and the risk of large-scale denial-of-service failures remain open concerns. Thus, while our architecture advances IoT edge security and performance, these domain-specific factors highlight the importance of real-world validation under practical deployment conditions.

4.4. Advantages over previous approaches

The integration of Lightweight Virtualization (LV) and Trusted Execution Environments (TEE) offers clear advantages compared to earlier approaches that relied on either technology in isolation. First, the combined architecture significantly improves performance metrics, achieving a 40.93% reduction in latency, a 19.95% increase in throughput, and a 33.65% reduction in energy consumption over existing frameworks. While previous TEE-based works provided strong isolation, they often introduced higher overhead (20–25%), and virtualization-only methods lacked robust hardware-enforced security. Our approach reduces overhead to 10–15% while maintaining real-time performance. Second, the dual-layer defense enhances security resilience, as LV ensures workload isolation while TEE provides hardware-backed protection for sensitive operations, thereby mitigating a wider spectrum of attacks, including side-channel exploits and malware injection. Third, the architecture demonstrates greater scalability and adaptability, reducing resource utilization to 60–70% compared to 75–85% in prior works, making it better suited for battery-powered IoT devices. Finally, this integration improves cost-effectiveness by lowering both resource consumption and energy usage, enabling sustainable deployment in large-scale IoT ecosystems.

4.5. Managerial implications

The findings of this research provide several important implications for managers, industry leaders, and policymakers overseeing IoT-enabled ecosystems:

(1) The proposed architecture offers a robust defense against cyber threats by integrating LV and TEEs. Managers in critical sectors such as healthcare, smart cities, and industrial automation can leverage this to build trust with stakeholders and strengthen compliance with data protection regulations.
(2) The architecture supports seamless expansion of IoT networks with minimal overhead. This provides managers with a cost-effective pathway to scale operations without compromising security or efficiency.

(3) Trust mechanisms such as secure boot and remote attestation enhance the organizationâ€™s ability to meet compliance requirements (e.g., GDPR, critical infrastructure standards). This reduces legal and reputational risks while improving resilience to cyberattacks.
(4) As edge computing becomes the foundation of next-generation IoT, managers can utilize these insights to prioritize investments in lightweight, secure, and scalable edge technologies that future-proof their infrastructure.

In essence, the contribution of this research extends beyond technical innovation by providing managers with actionable strategies to enhance security, reduce costs, comply with regulations, and achieve sustainable growth in IoT-driven industries.

5. Conclusions

In conclusion, the integration of lightweight virtualization (LV) and trusted execution environments (TEEs) significantly enhances the security and performance of IoT edge devices. The proposed method, which combines both LV and TEE, offers clear improvements in key metrics such as latency with 40.93% reduction, throughput with a 19.95% improvement, and energy consumption with a 33.65% reduction. As demonstrated, the combination of LV and TEE results in lower latency, higher throughput, and reduced energy consumption compared to using either technique individually. This makes the combined approach particularly well-suited for IoT edge devices that require real-time processing and efficient resource management. Edge computing, when supported by LV and TEE, enables IoT applications to operate with ultra-low latency and optimized security, positioning the edge layer as the backbone of future IoT ecosystems. By bringing cloud-based services closer to the network, edge computing not only improves mobility and location awareness but also strengthens the real-time performance of IoT applications. This research has provided a detailed comparison of existing cloud-based and edge-based paradigms, demonstrating how edge computing can address the limitations of traditional IoT systems, particularly when coupled with advanced virtualization and security solutions. While the integration of lightweight virtualization and TEEs represents a significant step forward, certain challenges remain. The need for further research in areas such as efficient resource management, dynamic task scheduling, and scalability is evident. Additionally, the ability to maintain high performance across increasingly complex and resource-constrained environments will be crucial as IoT networks continue to expand.

Acknowledgments

None.

Funding

The first three authors thank DST-FIST for funding the lab facility for supporting this research under grant number SR/FST/ET-II/2019/450.

Conflict of interest

The authors declare that they have no conflict of interest regarding the publication of this article.

Author contributions

Conceptualization: Ramakrishna Goli
Formal analysis: Ramakrishna Goli, Kumar Sureshkumar, Sundarakannan Mahilmaran
Methodology: Ramakrishna Goli, Aravindhan Alagarsamy, Gian Carlo Cardarilli
Writing–original draft: Ramakrishna Goli, Aravindhan Alagarsamy
Writing–review & editing: All authors

Availability of data

Not applicable.

AI tools statement

All authors confirm that no AI tools were used in the preparation of this manuscript.

References

1. Buyya R, Yeo CS, Venugopal S, Broberg J, Brandic I. Cloud computing and emerging IT platforms: vision, hype, and reality for delivering computing as the 5th utility. *Future Gener Comput Syst.* 2009;25(6):599-616.
2. Puliafito C, Mingozzi E, Longo F, Puliafito A, Rana O. Fog computing for the internet of things: a survey. *ACM Trans Internet Technol.* 2019;19(2):1-41.
3. Mansouri Y, Toosi AN, Buyya R. Data storage management in cloud environments: taxonomy, survey, and future directions. *ACM Comput Surv.* 2017;50(6):1-51.
4. Kocher P, Horn J, Fogh A, et al. Spectre attacks: exploiting speculative execution. *Commun ACM.* 2020;63(7):93-101.
5. Mosenia A, Jha NK. A comprehensive study of security of Internet of Things. *IEEE Trans Emerg Top Comput.* 2016;5(4):586-602

6. Yang Y, Wu L, Yin G, Li L, Zhao H. A survey on security and privacy issues in Internet of Things. *IEEE Internet Things J.* 2017;4(5):1250-1258
7. Zhang J, Chen B, Zhao Y, Cheng X, Hu F. Data security and privacy-preserving in edge computing paradigm: survey and open issues. *IEEE Access.* 2018;4:18209-18237
8. Zanella A, Bui N, Castellani A, Vangelista L, Zorzi M. Internet of Things for smart cities. *IEEE Internet Things J.* 2014;1(1):22-32
9. Atzori L, Iera A, Morabito G. The Internet of Things: a survey. *Comput Netw.* 2010;54(15):2787-2805
10. Gubbi J, Buyya R, Marusic S, Palaniswami M. Internet of Things (IoT): a vision, architectural elements, and future directions. *Future Gener Comput Syst.* 2013;29(7):1645-1660
11. Catarinucci L, De Donno D, Mainetti L, et al. An IoT-aware architecture for smart healthcare systems. *IEEE Internet Things J.* 2015;2(6):515-526.
12. Chen N, Chen Y. Smart city surveillance at the network edge in the era of IoT: opportunities and challenges. *Smart Cities: Dev Gov Frameworks.* 2018:153-176.
13. Junior FM, Kamienski CA. A survey on trustworthiness for the Internet of Things. *IEEE Access.* 2021;9:42493-42514.
14. Fitwi A, Chen Y, Zhu S, Blasch E, Chen G. Privacy-preserving surveillance as an edge service based on lightweight video protection schemes using face de-identification and window masking. *Electronics (Basel).* 2021;10(3):236.
15. Gisdakis S, Lagana M, Giannetsos T, Papadimitratos P. SEROSA: service-oriented security architecture for vehicular communications. In: *Proc IEEE Vehicular Networking Conference*; 2013:111-118.
16. Dimitriou T, Giannetsos T, Chen L. REWARDS: privacy-preserving rewarding and incentive schemes for the smart electricity grid and other loyalty systems. *Comput Commun.* 2019;137:1-4.
17. Wang W, Tornatore M, Zhao Y, et al. Infrastructure-efficient virtual machine placement and workload assignment in cooperative edge-cloud computing over backhaul networks. *IEEE Trans Cloud Comput.* 2021;11(1):653-665.
18. Liang J, Li K, Liu C, Li K. Joint offloading and scheduling decisions for DAG applications in mobile edge computing. *Neurocomputing.* 2021;424:160-171.
19. Al-Habob AA, Dobre OA, Armada AG, Muhaidat S. Task scheduling for mobile edge computing using genetic algorithm and conflict graphs. *IEEE Trans Veh Technol.* 2020;69(8):8805-8819.
20. Zhang F, Zhang H. SoK: a study of using hardware-assisted isolated execution environments for security. In: *Proc. of Hardware Archit Support Secur Priv.* 2016;3:1-8.
21. Al-Omary A, Othman A, AlSabbagh HM, Al-Rizzo H. Survey of hardware-based security support for IoT/CPS systems. *KnE Eng.* 2018:52-70.
22. Adams K, Agesen O. A comparison of software and hardware techniques for x86 virtualization. *ACM Sigplan Not.* 2006;41(11):2-13.
23. Khan MN, Rao A, Camtepe S. Lightweight cryptographic protocols for IoT-constrained devices: a survey. *IEEE Internet Things J.* 2020;8(6):4132-4156.
24. Madria S, Kumar V, Dalvi R. Sensor cloud: a cloud of virtual sensors. *IEEE Softw.* 2013;31(2):70-77.
25. Santos IL, Pirmez L, Delicato FC, et al. A resource allocation algorithm for the cloud of sensors. *Future Gener Comput Syst.* 2019;92:564-581.
26. Sahni Y, Cao J, Zhang S, Yang L. Edge mesh: a new paradigm to enable distributed intelligence in Internet of Things. *IEEE Access.* 2017;5:16441-16458.
27. Hoang TT, Duran C, Serrano R, et al. Trusted execution environment hardware by isolated heterogeneous architecture for key scheduling. *IEEE Access.* 2022;10:46014-46027.
28. Kumar VB, Chattopadhyay A, Haj-Yahya J, Mendelson A. Itus: a secure RISC-V system-on-chip. In: *Proc. International System-on-Chip Conference Sep 3*; 2019:418-423.
29. Haj-Yahya J, Wong MM, Pudi V, Bhasin S, Chattopadhyay A. Lightweight secure-boot architecture for RISC-V system-on-chip. In: *Proc International Symposium on Quality Electronic Design*; 2019:216-223.
30. Lee D, Kohlbrenner D, Shinde S, Asanović K, Song D. Keystone: an open framework for architecting trusted execution environments. In: *Proc of the Fifteenth European Conference on Computer Systems*; 2020:1-16.
31. Bahmani R, Brasser F, Dessouky G, et al. CURE: a security architecture with customizable and resilient enclaves. In: *Proc. of USENIX Security Symposium*; 2021:1073-1090.
32. Nasahl P, Schilling R, Werner M, Mangard S. HECTOR-V: a heterogeneous CPU architecture for a secure RISC-V execution environment. In: *Proc. of the 2021 ACM Asia Conference on Computer and Communications Security*; 2021:187-199.
33. Costan V, Lebedev I, Devadas S. Sanctum: minimal hardware extensions for strong software isolation. In: *Proc. of USENIX Security Symposium*; 2016:857-874.
34. Xia K, Luo Y, Xu X, Wei S. SGX-FPGA: trusted execution environment for CPU-FPGA heterogeneous architecture. In: *Proc. of ACM/IEEE Design Automation Conference*; 2021:301-306.
35. Cilardo A. Memory encryption support for an FPGA-based RISC-V implementation. In: *Proc. of International Conference on Design Technology of Integrated Systems in Nanoscale Era*; 2021:1-5.

36. Aitchison C, Buckle R, Ch'ng A, Clarke C, Malley J, Halak B. On the integration of physically unclonable functions into ARM TrustZone security technology. In: *Proc. of European Conference on Circuit Theory and Design*; 2020:1-4.
37. Armanuzzaman M, Zhao Z. BYOTee: towards building your own trusted execution environments using FPGA. *arXiv Preprint.* 2022;arXiv:2203.04214.
38. Meng X, Raj K, Ray S, Basu K. SeVNoC: security validation of system-on-chip designs with NoC fabrics. *IEEE Trans Comput Aided Des Integr Circuits Syst.* 2023;42(2):672-682.
39. Singh SK, Pan Y, Park JH. OTS scheme-based secure architecture for energy-efficient IoT in edge infrastructure. *Comput Mater Contin.* 2021;66:2905-2922.

Ramakrishna Goli *is a research scholar in Department of Electronics and communication Engineering associated with Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur Dist., AP, India. His research interest includes Optimization, Internet of Things, ASIC and FPGA*
https://orcid.org/0009-0003-6855-6283

Aravindhan Alagarsamy *received his B.E. in Electrical and Electronics Engineering from Madurai Kamaraj University, Madurai, India in 2003. He was awarded the M. Tech., Degree in VLSI Design from Kalasalingam Academy of Research and Education, Tamil Nadu India in 2009, and Ph.D., Degree from National Institute of Technology, Tiruchirappalli in the area of Networks-on-Chip. Currently, he is in the position of Associate Professor cum Associate Dean (Academics), Department of ECE, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur Dist., AP, India. He had industrial experience as Design Engineer in Electronic Design Automation Industry. His research interest includes Network on Chip, Optimization, ASIC, FPGA, and Soft Computing.*
https://orcid.org/0000-0003-3945-5080

Kumar Sureshkumar *received his B.E. in Electronics and Communication Engineering from IFET College of Engineering and Technology, Villupuram, Tamil Nadu, India, in 2007. He completed his M.E. in Communication Systems at Hindustan Institute of Technology, Chennai, Tamil Nadu, India, in 2011, and later obtained his Ph.D. from Annamalai University, Chidambaram, Tamil Nadu, India. He is currently serving as an Assistant Professor in the Department of Electronics and Communication Engineering at Koneru Lakshmaiah Education Foundation. His research interests include Sensor Networks, Digital Signal Processing, Digital Communication, and Wireless Communication.*
https://orcid.org/0000-0002-2421-4877

Sundarakannan Mahilmaran *received the B.Sc.(Mathe matics) degree in 2003 and the M.Sc. degree in Mathematics in 2005 from Madurai Kamaraj University, Madurai, India. He was awarded the Ph.D. degree in Mathematics at National Center for Advanced Research in Discrete Mathematics, Kalasalingam University, Tamil Nadu, India in 2011. His research interest includes discrete mathematics and theoretical computer science. Currently, he is working as an Assistant Professor (Grade 3), Department of Mathematics, Sri Sivasubramaniya Nadar(SSN) College of Engineering, Chennai, India since 2011.*
https://orcid.org/0000-0001-8902-1307

Gian Carlo Cardarilli *(Life Member, IEEE) was born in Rome, Italy. He received the Laurea degree (summa cum laude) from Sapienza UniversitÃ di Roma, in 1981. Since 1984, he has been with the Tor Vergata University of Rome, where he is currently a Full Professor of digital electronics and electronics for communication systems. From 1992 to 1994, he was with the University of Lâ€™Aquila. From 1987 to 1988, he was with the Circuits and Systems Team, EPFL, Lausanne, Switzerland. He works in the field of computer arithmetic and its application to the design of fast signal digital processors. He has also regular cooperation with companies, such as Alcatel Alenia Space, Italy; STM, Agrate Brianza, Italy; Micron, Italy; and Selex S.I., Italy. His research interests include VLSI architectures for signal processing and IC design. In this field, he has published more than 160 papers in international journals and conferences. His scientific interest includes the design of special architectures for signal processing.*
https://orcid.org/0000-0002-7444-876X

An International Journal of Optimization and Control: Theories & Applications
(https://accscience.com/journal/ijocta)