

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

Enhancing Security in Internet of Things (IoT) Systems through Data Provenance

Shahbaz Ali Mirjat

Department of Computer science, University of Sindh

shahbazalimirjat100@gmail.com

Muhammad Hamza Hussain

School of System and Technology ,University of Management and Technology Lahore

hamxzaa153@gmail.com

Amber Sarwar Hashmi

Department of computer science Rawalpindi Women university

amber.hashmi@rwu.edu.pk

Shoaib Hayat

Department of Computer Science and Technology Auckland University of Technology, (AUT) Auckland, New Zealand

shoaib.devpy@gmail.com

Abdul Hakeem

Department Computer Science Shaheed Zulfiqar Ali Bhutto Institute of Science and Technology (SZABIST), Islamabad

Email: hakeemabd638@gmail.com

ABSTRACT

The Internet of Things (IoT) is simply a collection of billions of resource-limited, device heterogeneity and constantly producing/delivering data on untrusted networks that perimeter-based security mechanisms are unable to defend. Many IoT security incidents are not because of a single point of failure, but because the data source is unknown, the data is altered, or whether it is trusted or not before it makes a decision. In this paper, the authors explore the concept of data provenance, defined as the systematic documentation of the origin, custody and evolution of data, as a key mechanism to improve IoT security. The paper examines the threat landscape of IoT and surveys various approaches to provenance capture and verification, such as blockchain-based ledgers, hardware-rooted trust anchors, cryptographic hash chaining, and machine-learning assisted provenance graph analysis, and introduces a five-layer provenance-enabled security architecture that combines lightweight cryptographic tagging at the perception layer, with distributed ledger commitment and anomaly-aware auditing at the application layer. Two summary tables provide a comparative analysis

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

of the techniques of provenance on overhead, tamper-resistance and scalability axes, and three explanatory figures help clarify the concepts. The benefits of security guarantees are discussed along with the restricted computational, memory and energy resources of IoT devices; open challenges such as standardisation, interoperability, proving it is not quantum-resistant, and disclosure with privacy are identified. Overall, the paper suggests that provenance-aware security, in conjunction with lightweight cryptography and edge-based verification, provides a viable and scalable approach to secure and trustworthy IoT ecosystems.

Keywords— Internet of Things, data provenance, IoT security, blockchain, trust management, lightweight cryptography, intrusion detection

1. Introduction

The Internet of Things (IoT) has become an engineering idea and paradigm that influences smart homes, industrial automation, intelligent healthcare, precision agriculture, and smart transportation systems. Today, billions of sensors and actuators, and embedded controllers, continuously monitor the physical world, communicate over heterogeneous networks, and drive automated decision pipelines with increasingly little or no human supervision. This connectivity on this scale has created an economic and operational value that is huge, but it's also created an attack surface that is both large and varied. In contrast to traditional information technology systems, IoT devices are often resource-constrained, physically vulnerable, with varying hardware and firmware capabilities, and may be placed for years with infrequent patching, making them easily targeted by adversaries to gain access to data integrity, cause disruption, or enter into larger enterprise networks. One thing that is unique about an IoT security incident is that the data is not just breached through unauthorized access but also from data corruption or falsification in the sensing, transmission, aggregation, and decision-making process.

A spoofed sensor, firmware update via a compromised gateway, or tampered log entry in an industrial control system can all cause incorrect automated actions before a traditional IDS can detect and alert. Traditional security approaches like perimeter firewalls, transport-layer encryption, or signature-based intrusion detection systems primarily focus on securing data while it's in transit, or securing that it's not accessed by someone else; they don't address the more basic question of whether the data is received from an origin and has been reliably held and properly manipulated by a system to maintain its integrity. Data provenance provides an answer to this question. Provenance is the history of a data element, including its origin and who has created, modified, or transmitted it, and how and when it has been transformed into its current form. While provenance has been the subject of research in scientific workflows and in the database field, in the last few years, it has gained traction in the field of IoT security, facilitated mainly by the advancement of blockchain and distributed ledger technologies, lightweight cryptographic primitives supported by constrained devices, and graph-based machine learning that is able to represent causal relationships among IoT events. The recent systematic literature reviews confirm that the study of security with provenance is an emerging and fast-growing research field, while several open

questions remain: scalability, disclosure with privacy, and unification of provenance models for different heterogeneous IoT ecosystems. In this paper, the authors explore systematic integration of data provenance into IoT security architectures, to boost the trustworthiness of data, facilitate forensic investigation and allow early detection of anomalous or malicious behavior. In particular, the paper aims at four goals. It first examines the IoT threat landscape and identifies how conventional security controls fail to address these threats, and how provenance is uniquely positioned to fill the gap. Second, it reviews the main families of provenance capture and verification methods from the literature, such as blockchain-based ledgers, hardware anchored trust, cryptographic hash chaining and provenance-graph-based intrusion detection. Third, it outlines a layered, provenance-enabled security architecture designed to suit the resource limitations of IoT deployments, as well as the lifecycle of the provenance records themselves and how they are created, committed, and verified. Fourth, it describes a comparison of the different types of provenance techniques for overhead, tamper-resistance, and scalability, and outlines the open challenges that need to be addressed before implementing provenance-based security at Internet scale. The rest of the paper is structured as follows. In Section 2, previous research in the field of IoT security and data provenance is surveyed. Section 3 describes the threat landscape of the IoT. Basic concepts and models of data provenance are introduced in Section 4. The proposed provenance-enabled security architecture is presented in Section 5. The methodology of structuring the review and comparative analysis is presented in Section 6. The comparative analysis is presented in Section 7 with the appropriate tables and figures. The findings are discussed in Section 8. Section 9 outlines open challenges and future research directions and Section 10 concludes this paper.

2. Literature Review

Data provenance research has been around for over a decade before the IoT age. Early scientific computing research established the concept of provenance as a directed graph of entities, activities and agents, and introduced models for the recording of the derivation of computational results, which were later standardised in the W3C PROV data model (Freire et al., 2008). Secure storage systems researchers around the same time showed that a cryptographically chained provenance log can thwart even if an adversary gets its hands on the storage medium itself (Hasan et al., 2009). These fundamental contributions laid the groundwork for provenance based security today, through a formal representation of data lineage and a mechanism for documenting data lineage in a tamper-proof way.

With the advent of lightweight, device-level provenance schemes that fit the low-resource sensor nodes, the idea of extending the concept of provenance to IoT systems started gaining traction. Aman et al. (2017) presented one of the first secure data provenance schemes explicitly developed for the IoT environment, where a secure data provenance binding is achieved by associating the data provenance with the device identity by leveraging the use of PUFs, without imposing high cryptographic overhead. Building on this work, Kamal and Tariq (2018) introduced a lightweight provenance scheme for multi-hop IoT networks, showing that provenance metadata can be carried with transmitted packets with minimal overhead in terms of bandwidth. These early

works were critical to demonstrating that lightweight security does not need to be incompatible with provenance, where cryptographic primitives are carefully chosen to be appropriate for the device capabilities.

The addition of blockchain technology has significantly transformed the future of research on the topic of the provenance. Some of the earlier, systematic arguments for blockchains to offer a tamper-resistant, decentralized ledger for IoT provenance are due to Christidis and Devetsikiotis (2016). Based on this assumption, later works have put forward solid solutions for blockchain-based IoT provenance frameworks that consist of specific functional requirements, including device diversity, ownership changes, and query efficiency. Recent systematic reviews have brought together this ever-increasing work: Faraj et al. (2025) performed a systematic literature review on security solutions for data provenance in IoT, and concluded that, although the majority of the literature focused on blockchain-based solutions, there is still a lack of research that evaluates these solutions in realistic, large-scale IoT workloads. In the same way, Jalbani, Weerawarna, and Al-Zubaidi (2025) suggested a quality model for provenance in AI-enhanced systems, highlighting that the guarantees of provenance which are now expected of the raw sensor data must now be extended to the machine learning pipelines consuming that data.

Provenance techniques have also been used in a similar way to identify and trace security events once data has been ingested by a system. Provenance-graph-based intrusion detection is a system that models system activity as a graph of processes, files, and network events, then looks for system events that are consistent with known or suspicious attack patterns. Representative examples of such systems include Unicorn (Han et al., 2020), which models long-running behavior of the system to detect advanced persistent threats, Shadewatcher (Zeng et al., 2022), which leverages graph learning via a recommendation approach to the threat analysis problem, and more recently presented systems such as Kairos (Cheng et al., 2024), Flash (Rehman, Ahmadi, & Hassan, 2024), and R-CAID (Goyal, Wang, & Bates, 2024) all presented at top-tier security conferences in 2024, and reported significant improvements in detection accuracy and root-cause analysis efficiency over previous approach using provenance-graphs. These techniques were extended to industrial and embedded network environments by Anjum, Iqbal, and Hamelin (2022) with their idea of a provenance-graph framework, specifically for advanced persistent threat detection. This body of work is discussed broadly by Pan et al. (2023) in the context of the broader intersection of data provenance, security and privacy.

In addition to work using graphs, there is a large body of research in machine learning and deep learning that focuses specifically on traffic and behavioral data from IoT devices and makes use of this data to detect anomalies to help identify potential violations of provenance, spoofed devices, or injected data. DeMedeiros, Hendawi, and Alvarez (2023) reviewed AI-based anomaly detection approaches for IoT and sensor networks, and cited the growing adoption of graph neural networks (GNN) to model the relationships among devices and traffic. Khan and Alkhathami (2024) showed that machine learning could be applied effectively for anomaly detection in an IoT-based healthcare system which is an application with a high provenance sensitivity, owing to the criticality of erroneous vital-sign data in healthcare. A deep-learning based anomaly

detection model was proposed by Alsoufi et al. (2024) and the studies using federated learning (Abbas et al., 2023) and ensemble deep neural networks (Thakkar & Lohiya, 2023) demonstrate the ongoing efforts to enhance detection accuracy and computational efficiency for constrained IoT deployment. These studies have been facilitated by the use of real benchmarks, like the CICIoT2023 (Neto et al., 2023) benchmark datasets.

Lastly, an increasing number of studies focus on lightweight cryptographic primitives that can enable the verification of provenance and integrity on constrained IoT devices. The National Institute of Standards and Technology completed its multi-year lightweight cryptography standardization project in February 2023, selecting the Ascan family of algorithms for authenticated encryption and hashing with significantly reduced computation and energy consumption compared to existing primitives like AES-GCM and SHA-256 which were previously too burdensome for embedding provenance tags into constrained sensor firmware. The other related works of lightweight signcryption (Kim et al., 2021) and secure aggregated data sharing in wireless sensor networks (Said et al., 2022) show that the objective of confidentiality, integrity and provenance can be met increasingly simultaneously rather than as additive security layers. Overall, the literature suggests a clear path from the early light-weight device-centric provenance systems to the distributed ledger systems based on blockchain to the existing graph-based and machine learning-enabled provenance analysis systems, with a common theme of lightweight cryptography. The present paper takes the thread of this development by combining these lines into a layered architecture and contrasting the major technique families on dimensions relevant for IoT deployment.

3. IoT Security Threat Landscape

Three structural characteristics of IoT deployments – device heterogeneity, resource constraints and physical exposure – create a threat landscape that differs from a traditional IT environment. The devices also come from several vendors whose security maturity levels vary, meaning that a single IoT deployment can include outdated device firmware, weak or hardcoded credentials, and devices that don't support common security protocols. Many devices, however, have limited computing power, memory and battery life and can't always support the computationally intensive cryptographic operations that security needs in an enterprise would require, forcing designers to make security compromises. Devices are often placed in physically accessible or unmonitored locations and can be vulnerable to tampering in the physical device, side-channel attack, and device cloning, which is not possible for data-center-hosted servers. In this context, threats can be structured by the lifecycle phase that they attack. When attacks are possible on the sensor data generation points, the data can be tampered with by an attacker presenting the same sensor as a different one with false readings that are statistically realistic, and therefore not detected by purely statistical analysis of the data, while a provenance record which is not able to verify the data against the sensor's registered identity would immediately identify the attack. While in transit, man-in-the-middle and replay attacks can modify or duplicate legitimate data as it passes through gateways and multi-hop mesh networks; hash-chained provenance records can detect those modifications even if the transport channel on which the data is sent is not fully

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

trusted. Compromised gateways or cloud back-ends can affect historical records in ways that are not detectable once the data is aggregated or stored, an issue append-only, distributed provenance ledgers can help prevent. Last but not least, where the data is used (the point of consumption), decision-making systems with no means for verifying data lineage can be susceptible to relying on data where the trustworthiness cannot be guaranteed, such as in safety-critical areas like connected vehicles, industrial control, and remote healthcare monitoring.

In addition to these threats to the life cycles of IoT, the literature and large-scale incident data reveal multiple classes of attacks that can be leveraged against the IoT ecosystem. While most likely used to test the integrity of the system, such attacks are also commonly used to deny service on a massive scale due to weak default credentials, typically provided by the Mirai botnet, and often as the first step to a more focused assault on the data when an initial foothold has been secured. Firmware attacks and supply-chain attacks introduce malicious code into the devices before or during deployment and modify the sensor readings—either by falsification or by data exfiltration—providing a kind of attack that is harder to detect with traditional signature-based methods. In IoT environments for industrial and critical infrastructures, the dwell time of an APThreat and the low number of events per minute have enabled it to evade detection through log-based surveillance systems: provenance-graph-based threat detection systems like Kairos and Flash have been proven to be able to offer meaningful improvements for detecting APThreats by using causal reasoning over events that is able to extend the time horizon of the monitoring.

Throughout this threat landscape we see that security controls are essential and must be configured to ensure confidentiality, but they are not enough when it comes to security. A sensor reading encrypted during transit has no direct association with the sensor's legitimacy or the absence of compromise or has no protection against an attacker with a valid key who can introduce bogus readings. This separation encourages the adoption of provenance as a complementary security goal, not an alternative one: Provenance mechanisms are specifically intended to answer questions of origin, custody and integrity history that can't be answered by access control and encryption; and they do not only prevent unauthorised access before it happens, but they also provide the forensic trail required to investigate an incident after it happens..

Table 1. IoT Security Threats Mapped to the Data Lifecycle and Corresponding Provenance-Based Countermeasures

Threat Category	Description	Lifecycle Stage	Provenance-Based Countermeasure
Sensor spoofing / device impersonation	Falsified or forged sensor readings injected by an unauthorized or cloned device	Data generation	Device-bound cryptographic identity; hardware-rooted signing of provenance tags
Man-in-the-middle / replay	Interception, alteration, or	Transmission	Hash-chained provenance records;

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

Threat Category	Description	Lifecycle Stage	Provenance-Based Countermeasure
attack	duplication of data in transit across gateways		sequence-numbered, timestamped tags
Retroactive log/data tampering	Silent modification of historical records at a compromised gateway or server	Aggregation / storage	Append-only distributed ledger commitment; periodic external anchoring
Firmware supply-chain compromise /	Malicious logic embedded in devices before or during deployment	Data generation	Provenance-graph anomaly detection; trust-score degradation on tag irregularity
Advanced persistent threat (low-and-slow)	Long-dwell, causally coordinated intrusion evading point-in-time detection	Aggregation / consumption	Provenance-graph representation learning across extended time horizons
Untrusted decision input	Downstream system acts on data of unverifiable origin or history	Consumption	Mandatory provenance/trust-score verification gate prior to automated action

4. Data Provenance: Concepts and Models

In the IoT world, data provenance is the metadata that can be used to trace back a data item's origin, the series of entities and processes that created it, modified it, or passed it along, and the timestamps and cryptographic evidence that allow that information to be duplicated and verified separately. The definition is based on the classical provenance model for scientific workflows, which associates entities, activities and agents with a directed acyclic graph, but adapted to the constraints and threat-model specific to IoT: the provenance records must be small enough to be sent down constrained links, cheap enough to be generated by a low-power microcontroller, and resistant to tampering by an adversary that might control the microcontroller (or even the gateway) that is generating the record.

There are typically three properties that are considered necessary for IoT provenance to be useful for securing the IoT network. The first is integrity, that is, the initial record of a transaction cannot be changed without anyone noticing it – this is usually enforced by cryptographic hash functions and digital signatures; sometimes these are linked together so that each new block contains a hash of the previous block, similar to a blockchain, and without a complete distributed ledger. The second is authenticity, that is, the ability to cryptographically attach a provenance record to the particular device

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

or entity that produced it, in a common way, using device bound keys, certificates, or physically unclonable functions which harness manufacturing-level variations to create a cryptographic identity for the device. The third is availability and queryability, that is, the availability of provenance records for verification and audit long after they are created raises practical questions about where to store provenance data, how long to store it, and the access-control policy governing it, especially when the availability of the data on many IoT devices is constrained by their storage resources.

The literature reports on a set of proven strategy for capturing the provenance that can be classified into four general families. Provenance records, or cryptographic digests of them, are stored on a blockchain with provenances of different networks distributed across a network of independent nodes, giving strong tamper-resistance and trust, but with high latency for transactions, rapidly increasing storage and for public blockchains, potentially large energy consumption. Hardware-rooted provenance uses trusted platform modules, secure elements, or physically unclonable functions in the device itself to generate and sign the provenance evidence and provides high-level guarantees of device authenticity, but has relatively high runtime costs and needs hardware support features that not all legacy IoT devices contain. The use of cryptographic hash-chaining provides a lightweight, software-only approach to retroactive tamper detection by making an append-only log, where the hash of each entry is embedded within the next entry, but this approach usually requires the existence of a single log-holder that must also be trusted or periodically anchored to an outside ledger. In the realm of intrusion detection, the most well-known use of provenances are those that are used to establish a graph of causal connections between events on a system including processes reading and writing files or network sockets, and then use graph analytics or graph neural networks to identify subgraphs that match with known patterns of malicious activity, using provenances not only to assert that a system has a lineage of data, but also to actively detect security violations in that lineage.

The four families are not colliding, but often they are mixed in proposed architectures—one is to have lightweight provenance evidence generated at the device level, using either hardware-rooted or hash-chained mechanisms, and periodically anchor digests of this evidence on a blockchain or consortium ledger for long-term tamper-evidence and cross-organizational trust; a separate analytics layer applies graph-based or machine-learning-based analysis to the accumulated evidence to detect anomalies. The architecture suggested in Section 5 of this paper is of this hybrid type and Section 7 compares the four technique families in terms of their overhead, tamper-resistance and scalability.

5. Proposed Provenance-Enabled Security Architecture

The proposed Provenance-Enabled Security Architecture is represented in Fig. 5. The proposed Provenance-Enabled Security Architecture is depicted in Fig. 5.

Drawing on the threat landscape and provenance concepts outlined earlier, this paper suggests a five-layer architecture with provenance for the security of IoT systems as shown in Figure 1. The design principle is that the provision of provenance evidence should be generated as near as possible to the point of creation of the data, propagated as part of the data, and verified independently at each stage of processing, to prevent

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

any compromise at any stage of the processing from silently corrupting the overall chain of trust.

The sensors, actuators, and embedded devices that collect raw data form the Perception Layer at the bottom of the architecture. The Perception Layer at the bottom of the architecture consists of sensors, actuators, and embedded devices that generate raw data. A lightweight cryptographic identity – preferably physically unclonable function or secure element – is provisioned into each device to enable it to sign a short provenance tag with each reading of data it sends. The provenance tag is intentionally sparse, consisting of a device identifier, a monotonic sequence number, a timestamp, and a small cryptographic "signature", computed using what is called a "lightweight primitive", which is a NIST-standardized primitive designed to produce authenticated encryption and hashing within the energy and memory constraints of devices capable of being used in microcontroller-based systems.

The Network and Edge Layer are composed of gateways and fog nodes that combine data from many devices in the perception layer, convert between constrained and conventional network protocols, and do an initial verification pass of incoming provenance tags prior to passing data upstream. It is a key layer in the architecture since there is generally sufficient computational power available to implement more costly cryptographic verification, and it is a natural place to test for malformed, lost or duplicated provenance tags potentially indicating a compromised device or replay attack. The Edge nodes batch verified provenance tags to efficiently commit to the upstream, otherwise committing each individual sensor reading to a distributed ledger would incur a high transaction overhead.

The Provenance Store layer is a Distributed Ledger that receives batched provenance commitments from edge nodes and records them in a tamper-evident structure, which the proposed architecture considers as a design choice, not a hard requirement: when deployments have strong latency and energy constraints they could implement a permissioned consortium blockchain or a distributed ledger based on a directed-acyclic-graph shared among trusted organizational stakeholders; when deployments prioritize minimal infrastructure they may implement a cryptographically hash-chained log instead, anchored to an external ledger periodically. In both cases, only compact digests (not raw sensor data) are added to ledger, with the raw data being stored in conventional off-chain storage, and referenced via cryptographic hash—a design pattern that is always described in the literature as a way to manage storage growth and maintain scalability.

The P&T layer is above the ledger and maintains trust scores for devices and entities, access control policies for who can query and disclose provenance records, and reconstructs the provenance graph and links related data items over time. Also in this layer is the use of graph-based anomaly detection, as in the case of a provenance-graph based intrusion detection system, where the accumulated provenance history is not merely a passive audit trail but an active input into the ongoing security monitoring process, and can flag causally suspicious patterns, such as a device whose provenance tags suddenly come from an unexpected network path, well before a downstream application would otherwise detect anomalous behavior.

Lastly, the Application Layer receives verified data and a queryable summary of its

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

provenance, enabling downstream services, dashboards, and automated decision logic to decide whether or not to take any action based on the trustworthiness of the data received, such as to require a certain level of trust for the smart-grid controller to take action, or to flag as well as data with an incomplete or broken provenance chain for human input, but not for automated action. The corresponding end-to-end lifecycle a single data item follows through this architecture is illustrated in figure 2, starting with the data generation and metadata tagging, progressing through cryptographic signing and commitment on the ledger, verification and ongoing audits, and a feedback loop to influence trust scores for future metadata tagging.

Figure 1. Layered Architecture for Provenance-Enabled IoT Security




Figure 1. Layered architecture for provenance-enabled IoT security, showing the flow of data and provenance evidence from the perception layer to the application layer.

Figure 2. End-to-End Data Provenance Lifecycle in an IoT Pipeline




Figure 2. End-to-end data provenance lifecycle in an IoT pipeline, with a continuous feedback loop from audit to trust scoring.

6. Methodology

This paper uses a structured narrative review with a qualitative comparative analysis. The review process started by searching for targeted and archival papers published from the last five years (2020–2025) across major venues and databases, including IEEE Xplore, ACM Digital Library, Springer Link, and arXiv preprints from the most important security conferences, such as IEEE Symposium on Security and Privacy. Whenever possible, the literature was chosen to include systematic literature reviews and comprehensive surveys which give a broad overview of the literature, followed by technical papers which had an individual influence and introduced new mechanisms for provenance or new detection systems for provenance.

For each of the technique families introduced in Section 4, four different types of sources were organized: blockchain-based provenance, hardware-rooted and lightweight cryptographic provenance, provenance-graph-based intrusion detection, and machine-learning-based anomaly detection for IoT data. For every category, the review identified the security goal, the approach to the goal, and (if reported) reported numbers of computational overhead, detection accuracy, or scalability. As the reviewed studies vary in both the evaluation methods, datasets and hardware platforms, the comparison presented in Section 7 is a qualitative and relative comparison of the reported overhead and tamper-resistance trends, rather than a single controlled experimental benchmark; this is evident from the illustrative comparison of the overheads and tamper resistance trends in Figure 3, which is specifically labeled as a qualitative and relative comparison of the trends observed across the reviewed studies. To identify these architectural components and patterns of operation, the authors investigated the literature reviewed in Section 3, and to map these recurring patterns to the traditional layered architecture of IoT systems - perception layer, network layer and application layer - the authors explicitly mapped these recurring patterns to this architecture, extending it with a separate provenance and trust management layer and a distributed ledger layer that accounts for the particular needs of security based on provenance.

7. Comparative Analysis

Tables 2 and 3 and Figure 3 provide the integrated advantages that have been discussed in Sections 2 and 4 as well as representative studies. Tables 2 and 3 compare the 4 main families of techniques in terms of tamper resistance, computational overhead and scalability, and highlight representative studies from each family, respectively. An illustration of the relative overhead and tamper-resistance trend (qualitative) obtained from the literature reviewed is presented in figure 3.

Table 2. Comparative Attributes of Data Provenance Techniques for IoT Security

Technique	Tamper-Resistance	Computational Overhead	Scalability	Representative Studies
Public blockchain	Very high	High	Low (throughput-	Christidis & Devetsikiotis

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

Technique	Tamper-Resistance	Computational Overhead	Scalability	Representative Studies
			limited)	(2016); Faraj et al. (2025)
Consortium / permissioned blockchain	High	Moderate	Moderate	Novo (2018); Jalbani et al. (2025)
DAG-based distributed ledger	High	Moderate-low	Moderate- high	Faraj et al. (2025)
Hardware- rooted (TPM / PUF)	High (device- level)	Low	High	Aman et al. (2017); Kamal & Tariq (2018)
Cryptographic hash-chaining	Moderate	Very low	High	Hasan et al. (2009); Kim et al. (2021)
Provenance- graph analytics / ML	N/A (detection layer)	Moderate (analysis-side)	High (with sampling)	Han et al. (2020); Cheng et al. (2024); Rehman et al. (2024)

Table 3. Summary of Representative Provenance-Aware IoT Security Studies

Study (Year)	Technique Family	Primary Focus	Key Contribution
Aman et al. (2017)	Hardware-rooted provenance	Device-level authenticity for constrained IoT	PUF-based binding of provenance to device identity
Christidis & Devetsikiotis (2016)	Blockchain-based provenance	Decentralized trust for IoT data and transactions	Foundational case for smart-contract-managed IoT provenance
Han et al. (2020) – Unicorn	Provenance-graph detection	Long-running APT detection	Whole-system provenance modeling for stealthy intrusions
Zeng et al. (2022) – Shadewatcher	Provenance-graph detection	Recommendation-guided threat analysis	Graph-learning-based cyber threat analysis from audit records

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

Study (Year)	Technique Family	Primary Focus	Key Contribution
Cheng et al. (2024) – Kairos	Provenance-graph detection	Practical intrusion and investigation	Whole-system provenance at production-representative scale
Rehman et al. (2024) – Flash	Provenance-graph detection	Comprehensive intrusion detection	Representation learning directly on provenance graphs
Faraj et al. (2025)	Systematic literature review	State of IoT data provenance security research	Identifies evaluation and scalability gaps across the field
DeMedeiros et al. (2023)	ML-based anomaly detection	AI-based anomaly detection in IoT/sensor networks	Comprehensive survey of ML and DL detection approaches

Figure 3. Comparative Overhead and Tamper-Resistance of Data Provenance Techniques for IoT (Illustrative, based on reviewed literature)




Figure 3. Comparative overhead and tamper-resistance of data provenance techniques for IoT, synthesized from trends reported in the reviewed literature.

8. Discussion

Table 2 and Figure 3 show a consistent trade-off across the different approaches to provenance: approaches that provide the highest level of tamper-resistance tend to require the highest level of computation and communication, especially for public blockchains, and the approaches that require the least amount of computation and communication tend to provide the lowest level of tamper-resistance, typically relying on a single log-holder instead of a distributed set of validators. The implications of this trade-off for the architecture design are evident: a single provenance mechanism that operates at the global level of an entire IoT deployment will probably not be optimal

because each layer of the proposed architecture from Section 5 has different resource limitations and different consequences of a compromise. The distributed ledger layer, where aggregated data from numerous devices is committed, and tamper-resistance is crucial for long-term forensic and audit purposes, is better suited to a lightweight mechanism such as a consortium blockchain or similar distributed ledger, whereas the perception layer, which is where devices are at their most constrained and where each device is individually least consequential if compromised, is better served by lightweight hardware-rooted or hash-chained mechanisms.

Another learning from the reviewed literature is that of the convergence of provenance as a passive audit mechanism and provenance as an active detection mechanism. While early provenance models were used to validate a provenance record after an incident was suspected, models like Kairos, Flash, and R-CAID show that the provenance graph can be analyzed and continuously used as a signal for detecting causally coherent attack patterns that would not be detectable by event-oriented detection models. This convergence helps to align the architectural approach of including a continuous feedback loop from the audit step back into the trust scoring process, instead of a provenance verification gate on data consumption.

The third observation is on the role of machine learning as an augmenting layer to, not a substitute for, cryptographic provenance mechanisms. The machine-learning-based anomaly detection systems surveyed in Section 2 are good at determining if a traffic pattern or sensor measurement is statistically abnormal, but can usually not by themselves determine if the traffic pattern or sensor measurement is abnormal or is a cryptographically unauthenticated one, as they can look very similar. Cryptographic provenance mechanisms, on the other hand, can successfully determine if a data item's claimed origin is genuine, but cannot provide any certainty that an authenticated device is not compromised and is producing malicious, but validly signed data. This paper explores the architecture proposed where these are complementary layers of defense, with the cryptographic level of verifiable authenticity, and the machine learning analysis of the resulting provenance graph affording additional security on top of cryptographic security.

Lastly, the literature reviewed offers a consistent view of scalability as the biggest practical challenge in implementing provenance-based security in the scale commonly found in real world IoT environments, where millions of devices can be continuously producing data. While this is a challenge, committing compact digests instead of full records to the distributed ledger and batching up commitments in the edge layer instead of committing each individual read reduces, but does not eliminate, this challenge, and Section 9 discusses this and related open challenges in detail.

9. Challenges and Future Directions

Although significant advances are made, there are still some challenges open in the study and implementation of the provenance-based security for IoT. There are two: standardization and historical precedent. The W3C PROV model offers a widely accepted basis for provenance in scientific and enterprise computing, while no such widely accepted standard exists for IoT provenance metadata, which can make it difficult to leverage provenance evidence across different vendor and research-

developing provenance systems and can hinder interoperability across organizations for many industrial IoT and supply-chain use cases.

The second challenge is scaling up in realistic deployment scenarios. The literature reviewed indicates that provenance systems can achieve promising results during evaluation on testbeds of tens to hundreds of devices, but such tests on the scale of millions of devices concurrently operating, as would be expected in large-scale deployments for smart cities or industry, are less common. There remain questions about the throughput of ledgers at scale, how much storage they can accommodate as the number of devices grows, and how long it takes to run queries at Internet of Things scale.

The third challenge is disclosure through privacy-preserving disclosure. Provenance records are designed to provide information about the history and context of data and this information can be a privacy risk, especially in areas like connected healthcare where the detailed provenance record may provide information about a patient's location, activity or treatment history, even if the information contained in the sensor data is encrypted. The issue of reconciling transparency needed for provenance verification with the confidence needed in privacy regulations still has some research space to be explored and answered; selective disclosure and zero-knowledge proof methods have been pointed out in the literature as possible solutions, but are not widely used.

The fourth challenge is Post Quantum readiness. Most existing schemes are based on elliptic-curve or RSA-based digital signatures which could be broken by a quantum capable adversary, and the goal of these schemes is to be verifiable for years, so that the existence of quantum capable adversaries may be many years in the future and this presents an incentive for research into new post-quantum cryptographic schemes, particularly lattice-based and other schemes specifically designed for constrained edge devices.

The fifth challenge is economic and organizational, not just technical: provenance-based security needs coordinated investment from device makers, network operators and application developers, none of whom can capture the full security benefit from adopting provenance, resulting in another coordination problem much like other Internet-scale security transitions. It will probably take a mix of regulatory incentives, industry standardisation measures, and the proof of impact from decreased incident response and forensic investigation costs to overcome this challenge, a space that is under-explored in the current literature and is an exciting arena for further interdisciplinary research between technical, economic and policy perspectives.

10. Conclusion

This paper has studied data provenance as an underlying technology to improve security of IoT systems, given that traditional security approaches based on confidentiality and access control do not, on their own, provide trust in data prior to its use in automated decision making. This paper outlined a five-layer provenance-enabled security architecture, starting with generation of lightweight provenance evidence at the point of data creation and progressing through progressively more capable layers of security, from provenance evidence ledgers, to hardware-based trust anchors, to cryptographic

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

hash chaining, and to provenance-graph-based intrusion detection.

The comparative analysis in this paper shows that there is no one-size-fits-all provenance that is optimal for all IoT devices, but rather that creative IoT security architectures should intelligently utilize lightweight, device-level, tamper-resistance mechanisms that are suitable for the constraints of the perception-layer hardware, and selectively apply higher-overhead, higher-assurance distributed ledger mechanisms at the network and application layers where they are most valuable. The trend of recent literature of combining both passive, audit-type evidence and active input for provenance-based, machine-learning-based anomaly detection also indicates that the future of IoT security architectures is not an either-or situation, but in designing systems where both methods complement each other.

However, there are still some major hurdles to overcome before provenance-based security can be implemented at the Internet-of-Things scale: there is no widely accepted provenance metadata standard for IoT; tensions between provision of transparency and data privacy; requirements for post-quantum resistant provenance mechanisms; and organizational coordination necessary in a fragmented device manufacturing ecosystem. These challenges form a large but manageable research program, which this paper argues should be considered as a fundamental design principle of trustworthy IoT systems from an early stage of the device and architecture design. This paper has studied data provenance as an underlying technology to improve security of IoT systems, given that traditional security approaches based on confidentiality and access control do not, on their own, provide trust in data prior to its use in automated decision making. This paper outlined a five-layer provenance-enabled security architecture, starting with generation of lightweight provenance evidence at the point of data creation and progressing through progressively more capable layers of security, from provenance evidence ledgers, to hardware-based trust anchors, to cryptographic hash chaining, and to provenance-graph-based intrusion detection.

The comparative analysis in this paper shows that there is no one-size-fits-all provenance that is optimal for all IoT devices, but rather that creative IoT security architectures should intelligently utilize lightweight, device-level, tamper-resistance mechanisms that are suitable for the constraints of the perception-layer hardware, and selectively apply higher-overhead, higher-assurance distributed ledger mechanisms at the network and application layers where they are most valuable. The trend of recent literature of combining both passive, audit-type evidence and active input for provenance-based, machine-learning-based anomaly detection also indicates that the future of IoT security architectures is not an either-or situation, but in designing systems where both methods complement each other.

However, there are still some major hurdles to overcome before provenance-based security can be implemented at the Internet-of-Things scale: there is no widely accepted provenance metadata standard for IoT; tensions between provision of transparency and data privacy; requirements for post-quantum resistant provenance mechanisms; and organizational coordination necessary in a fragmented device manufacturing ecosystem. These challenges form a large but manageable research program, which this paper argues should be considered as a fundamental design principle of trustworthy IoT systems from an early stage of the device and architecture design.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

References

- Abbas, S., Sultan, S., Ghazal, R., Fatima, A., Alsubai, S., & Binbusayyis, A. (2023). A novel federated edge learning approach for detecting cyberattacks in IoT infrastructures. *IEEE Access*, 11, 112189–112198. <https://doi.org/10.1109/ACCESS.2023.3318866>
- Akbarfam, A. J., & Maleki, H. (2024). SoK: Blockchain for provenance. In *Proceedings of the VLDB 2024 Workshops: Fifth International Workshop on Applied AI for Database Systems and Applications (FAB)*. VLDB Endowment.
- Alabdulatif, A. (2022). Data provenance for cloud forensic investigations, security, challenges, solutions and future perspectives: A survey. *Journal of King Saud University – Computer and Information Sciences*, 34(10), 8330–8349. <https://doi.org/10.1016/j.jksuci.2022.08.006>
- AlJabri, M., Alahmadi, A. A., Mohammad, R. M. A., Alhamed, F., Altamimi, H. S., & Chrouf, S. M. B. (2023). A comprehensive review of lightweight authenticated encryption for IoT devices. *Wireless Communications and Mobile Computing*, 2023, Article 9071969. <https://doi.org/10.1155/2023/9071969>
- Alsoufi, M. A., Siraj, M. M., Ghaleb, F. A., Al-Razgan, M., Al-Asaly, M. S., Alfakih, T., & Saeed, F. (2024). Anomaly-based intrusion detection model using deep learning for IoT networks. *Computer Modeling in Engineering & Sciences*, 141(1), 823–845. <https://doi.org/10.32604/cmescs.2024.052112>
- Aman, M. N., Chua, K. C., & Sikdar, B. (2017). Secure data provenance for the internet of things. In *Proceedings of the 3rd ACM International Workshop on IoT Privacy, Trust, and Security* (pp. 11–14). Association for Computing Machinery. <https://doi.org/10.1145/3055245.3055255>
- Anjum, M. M., Iqbal, S., & Hamelin, B. (2022). Anubis: A provenance graph-based framework for advanced persistent threat detection. In *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing* (pp. 1684–1693). Association for Computing Machinery. <https://doi.org/10.1145/3477314.3507097>
- Bertino, E. (2014). Data trustworthiness—approaches and research challenges. In J. Garcia-Alfaro, G. Navarro-Arribas, N. Cuppens-Boulahia, & S. de Capitani di Vimercati (Eds.), *Data privacy management, autonomous spontaneous security, and security assurance* (pp. 17–25). Springer. https://doi.org/10.1007/978-3-319-17016-9_2
- Chaudhry, S. A. (2022). Comments on "A secure, privacy-preserving, and lightweight authentication scheme for VANETs". *IEEE Sensors Journal*, 22(13), 13763–13766. <https://doi.org/10.1109/JSEN.2022.3168512>
- Cheng, Z., Lv, Q., Liang, J., Wang, Y., Sun, D., Pasquier, T., & Han, X. (2024). Kairos: Practical intrusion detection and investigation using whole-system provenance. In *2024 IEEE Symposium on Security and Privacy (SP)* (pp. 3533–3551). IEEE. <https://doi.org/10.1109/SP54263.2024.00005>
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the internet of things. *IEEE Access*, 4, 2292–2303. <https://doi.org/10.1109/ACCESS.2016.2566339>

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

- DeMedeiros, K., Hendawi, A., & Alvarez, M. (2023). A survey of AI-based anomaly detection in IoT and sensor networks. *Sensors*, 23(3), Article 1352. <https://doi.org/10.3390/s23031352>
- Dorri, A., Kanhere, S. S., Jurdak, R., & Gauravaram, P. (2017). Blockchain for IoT security and privacy: The case study of a smart home. In 2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops) (pp. 618–623). IEEE. <https://doi.org/10.1109/PERCOMW.2017.7917634>
- Faraj, O., Megías, D., & Garcia-Alfaro, J. (2025). Security approaches for data provenance in the Internet of Things: A systematic literature review. *ACM Computing Surveys*, 57(10), 1–41. <https://doi.org/10.1145/3718735>
- Freire, J., Koop, D., Santos, E., & Silva, C. T. (2008). Provenance for computational tasks: A survey. *Computing in Science & Engineering*, 10(3), 11–21. <https://doi.org/10.1109/MCSE.2008.79>
- Goyal, A., Wang, G., & Bates, A. (2024). R-CAID: Embedding root cause analysis within provenance-based intrusion detection. In 2024 IEEE Symposium on Security and Privacy (SP) (pp. 3515–3532). IEEE. <https://doi.org/10.1109/SP54263.2024.00253>
- Han, X., Pasquier, T. F. J., Bates, A., Mickens, J., & Seltzer, M. I. (2020). Unicorn: Runtime provenance-based detector for advanced persistent threats. *arXiv*. <https://doi.org/10.48550/arXiv.2001.01525>
- Hasan, R., Sion, R., & Winslett, M. (2009). Preventing history forgery with secure provenance. *ACM Transactions on Storage*, 5(4), Article 12. <https://doi.org/10.1145/1629080.1629082>
- Henry, T., & Tucci-Piergiovanni, S. (2024). Scientific workflows management with blockchain: A survey. In *Blockchain and smart-contract technologies for innovative applications* (pp. 45–72). Springer.
- Hoopes, R., Hardy, H., Long, M., & Dagher, G. G. (2022). SciLedger: A blockchain-based scientific workflow provenance and data sharing platform. In 2022 IEEE 8th International Conference on Collaboration and Internet Computing (CIC) (pp. 1–10). IEEE.
- Jalbani, A., Weerawarna, R., & Al-Zubaidi, K. (2025). Enhancing data provenance in AI with blockchain technology: A comprehensive quality model. *CSI Transactions on ICT*. Advance online publication. <https://doi.org/10.1007/s40012-025-00419-7>
- Kamal, M., & Tariq, S. (2018). Light-weight security and data provenance for multi-hop internet of things. *IEEE Access*, 6, 34439–34448. <https://doi.org/10.1109/ACCESS.2018.2850821>
- Khan, K. M., Haider, W., Khan, N. A., & Saleem, D. (2023). Big data provenance using blockchain for qualitative analytics via machine learning. *Journal of Universal Computer Science*, 29(5), 446–468.
- Khan, M. A., & Salah, K. (2018). IoT security: Review, blockchain solutions, and open challenges. *Future Generation Computer Systems*, 82, 395–411. <https://doi.org/10.1016/j.future.2017.11.022>
- Khan, M. M., & Alkhathami, M. (2024). Anomaly detection in IoT-based healthcare:

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

- Machine learning for enhanced security. *Scientific Reports*, 14, Article 5872. <https://doi.org/10.1038/s41598-024-56126-x>
- Kim, T.-H., Kumar, G., Saha, R., Buchanan, W. J., Devgun, T., & Thomas, R. (2021). LiSP-XK: Extended light-weight signcryption for IoT in resource-constrained environments. *IEEE Access*, 9, 100972–100980. <https://doi.org/10.1109/ACCESS.2021.3097267>
- National Institute of Standards and Technology. (2023). Lightweight cryptography standardization process: NIST selects Ascon. U.S. Department of Commerce.
- Neto, E. C. P., Dadkhah, S., Ferreira, R., Zohourian, A., Lu, R., & Ghorbani, A. A. (2023). CICIOT2023: A real-time dataset and benchmark for large-scale attacks in IoT environment. *Sensors*, 23(13), Article 5941. <https://doi.org/10.3390/s23135941>
- Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2023). Federated learning for smart healthcare: A survey. *ACM Computing Surveys*, 55(3), Article 60. <https://doi.org/10.1145/3501296>
- Novo, O. (2018). Blockchain meets IoT: An architecture for scalable access management in IoT. *IEEE Internet of Things Journal*, 5(2), 1184–1195. <https://doi.org/10.1109/JIOT.2018.2812239>
- Pan, B., Stakhanova, N., & Ray, S. (2023). Data provenance in security and privacy. *ACM Computing Surveys*, 55(14s), Article 314. <https://doi.org/10.1145/3593294>
- Rehman, M. U., Ahmadi, H., & Hassan, W. U. (2024). Flash: A comprehensive approach to intrusion detection via provenance graph representation learning. In 2024 IEEE Symposium on Security and Privacy (SP) (pp. 3552–3570). IEEE. <https://doi.org/10.1109/SP54263.2024.00139>
- Restuccia, F., D'Oro, S., & Melodia, T. (2018). Securing the internet of things in the age of machine learning and software-defined networking. *IEEE Internet of Things Journal*, 5(6), 4829–4842. <https://doi.org/10.1109/JIOT.2018.2846040>
- Said, G., Ghani, A., Ullah, A., Azeem, M., Bilal, M., & Kwak, K. S. (2022). Light-weight secure aggregated data sharing in IoT-enabled wireless sensor networks. *IEEE Access*, 10, 33571–33585. <https://doi.org/10.1109/ACCESS.2022.3160231>
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646. <https://doi.org/10.1109/JIOT.2016.2579198>
- Shi, Y., Wei, W., Fan, H., Au, M. H., & Luo, X. (2019). A light-weight white-box encryption scheme for securing distributed embedded devices. *IEEE Transactions on Computers*, 68(10), 1411–1427. <https://doi.org/10.1109/TC.2019.2907847>
- Thakkar, A., & Lohiya, R. (2023). Attack classification of imbalanced intrusion data for IoT network using ensemble-learning-based deep neural network. *IEEE Internet of Things Journal*, 10(13), 11888–11895. <https://doi.org/10.1109/JIOT.2023.3244810>
- Wang, K. (2020). Security and privacy for edge intelligence in IoT: Challenges and solutions. *IEEE Network*, 34(5), 36–41.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

<https://doi.org/10.1109/MNET.001.2000050>

Wang, Z., Chen, H., Zhang, R., & Wang, Y. (2023). A lightweight intrusion detection method for IoT based on deep learning and dynamic quantization. *PeerJ Computer Science*, 9, Article e1569. <https://doi.org/10.7717/peerj-cs.1569>

Zeng, J., Wang, X., Liu, J., Chen, Y., Liang, Z., Chua, T.-S., & Chua, Z. L. (2022). Shadewatcher: Recommendation-guided cyber threat analysis using system audit records. In *2022 IEEE Symposium on Security and Privacy (SP)* (pp. 489–506). IEEE. <https://doi.org/10.1109/SP46214.2022.9833669>