International Journal of Computing and Artificial Intelligence

E-ISSN: 2707-658X
P-ISSN: 2707-6571
IJCAI 2021; 2(2): 70-79
Received: 02-12-2021
Accepted: 25-12-2021

Daniel J Agrinya
B.Eng. Mechanical Engineering, Cross River University of Technology, Calabar, Nigeria

Corresponding Author:
Daniel J Agrinya
B.Eng. Mechanical Engineering, Cross River University of Technology, Calabar, Nigeria

Improving Real-Time Threat Detection Accuracy Using Correlated Firewall, IDS, and Endpoint Telemetry Data

Daniel J Agrinya

DOI: https://www.doi.org/10.33545/27076571.2021.v2.i2a.258

Abstract
The growing complexity of modern enterprise networks has significantly increased the volume and diversity of security telemetry generated by firewalls, intrusion detection systems (IDS), and endpoint protection platforms. While these tools individually provide valuable visibility, reliance on isolated alerts often results in high false-positive rates, delayed detection, and analyst fatigue. From a broader cybersecurity perspective, improving threat detection accuracy requires integrated analysis that correlates signals across multiple control layers rather than evaluating events in isolation. This study narrows its focus to enhancing real-time threat detection accuracy through the correlation of firewall logs, IDS alerts, and endpoint telemetry data. It examines how multi-source data fusion, temporal alignment, and contextual enrichment can reveal attack patterns that are otherwise indistinguishable within single data streams. By correlating network-level activity with endpoint behavior and intrusion signatures, the proposed approach improves detection fidelity for lateral movement, credential abuse, and stealthy persistence techniques. The framework supports real-time analytics within security operations environments, enabling faster triage, reduced false positives, and more confident incident response decisions. The findings demonstrate that correlated telemetry analysis offers a scalable and practical method for strengthening threat detection capabilities in complex enterprise and hybrid network environments.

Keywords: Threat detection, Security telemetry, Firewall logs, Intrusion detection systems, Endpoint security, Event correlation

1. Introduction

1.1 Escalating Complexity of Enterprise Cyber Threats

Enterprise cyber threats have increased markedly in complexity due to the expansion of digital infrastructures across on-premise networks, distributed endpoints, and cloud-based platforms [1]. Modern enterprises operate hybrid environments where traditional perimeter boundaries are blurred, creating multiple entry points that adversaries can exploit simultaneously. This expansion of attack surfaces has been accompanied by the rise of stealthy, low-noise attack techniques designed to evade detection rather than trigger overt security alerts [3]. Adversaries increasingly favor credential abuse, living-off-the-land techniques, and slow lateral movement, allowing them to persist within environments for extended periods without raising suspicion. These tactics challenge conventional assumptions underlying many security controls, which are optimized to detect known patterns of malicious behavior or high-volume anomalies [5]. As attackers deliberately operate below detection thresholds, isolated controls struggle to distinguish malicious activity from legitimate system behavior. The result is a growing disparity between attacker sophistication and defensive visibility. This transition exposes a fundamental limitation: security mechanisms deployed in isolation lack the contextual awareness required to identify subtle, multi-stage intrusions that unfold across time and infrastructure layers. Consequently, enterprise security effectiveness increasingly depends not on the strength of individual controls, but on the ability to interpret complex activity patterns spanning diverse systems and environments [7].

1.2 Limitations of Single-Source Security Monitoring

Single-source security monitoring remains prevalent in many organizations, yet it provides only partial visibility into adversary behavior. Firewall-centric monitoring, for example,

primarily captures ingress and egress traffic at defined network boundaries, offering limited insight into internal lateral movement or malicious activity originating from trusted zones [2]. As enterprise architectures decentralize, reliance on perimeter-based telemetry alone obscures significant portions of the attack lifecycle.

Intrusion Detection Systems (IDS) introduce additional challenges. Signature-based IDS rely on predefined patterns of known attacks, rendering them vulnerable to evasion through polymorphism, encryption, or novel techniques [4]. Even anomaly-based systems often generate excessive alerts due to benign deviations in normal behavior, overwhelming analysts and diluting attention. This alert overload leads to desensitization, where critical signals are lost amid large volumes of low-confidence notifications [6].

Endpoint detection tools provide granular visibility into host-level activity, such as process execution and file modification, yet they frequently lack sufficient network or cross-host context to infer attacker intent [8]. Without correlation to network flows, authentication logs, or cloud activity, endpoint telemetry may appear innocuous in isolation. These fragmented monitoring approaches create a detection gap in which adversarial actions are distributed across multiple systems, each individually insufficient to trigger alarms. The absence of integrated visibility prevents security teams from reconstructing attack narratives, limiting their ability to detect coordinated or low-and-slow intrusions effectively [5].

1.3 Motivation for Correlated Telemetry Analysis

The limitations of isolated monitoring mechanisms motivate the need for correlated telemetry analysis that integrates signals across multiple security layers [1]. Correlation enables multi-layer threat visibility by combining network, endpoint, identity, and application data into a unified analytical context. Rather than evaluating events independently, correlated analysis identifies relationships and sequences that reveal attacker behavior patterns spanning diverse systems [3].

This approach shifts detection from retrospective investigation toward near real-time inference, allowing security teams to recognize evolving threats before significant impact occurs [7]. Correlation reframes intrusion detection as a detection-accuracy problem, where the objective is not to increase alert volume but to improve confidence through contextual aggregation. By aligning disparate telemetry streams, organizations can reduce false positives, surface stealthy attacks, and close visibility gaps inherent in single-source monitoring. Correlated telemetry analysis thus represents a foundational capability for modern enterprise threat detection, supporting adaptive defense in increasingly complex and adversarial environments [8].

2. Threat Detection Landscape and Telemetry Sources

2.1 Firewall Telemetry Characteristics and Constraints

Firewall telemetry represents one of the most established sources of security data within enterprise environments, primarily capturing network flow records and access control decisions at defined traffic inspection points [5]. Typical firewall logs include metadata such as source and destination addresses, ports, protocols, session durations, and allow or deny outcomes based on predefined rule sets. These records are valuable for establishing baseline traffic patterns, enforcing segmentation policies, and identifying overt policy violations. When analyzed longitudinally, firewall telemetry can reveal shifts in communication behavior that may indicate reconnaissance or data exfiltration attempts [7].

Despite this utility, firewall telemetry exhibits structural blind spots that limit its effectiveness for detecting modern threats. The widespread adoption of encrypted communication protocols significantly reduces visibility into payload content, constraining analysis to coarse-grained metadata [9]. While metadata may signal anomalies, it rarely provides sufficient context to distinguish malicious intent from legitimate encrypted services. Additionally, firewalls are typically positioned at network boundaries, resulting in limited insight into east-west traffic and internal lateral movement once an attacker gains an initial foothold. Internal communications between compromised hosts may bypass perimeter inspection entirely, remaining invisible to firewall-based monitoring [6].

These constraints mean that firewall telemetry alone tends to capture only fragments of the attack lifecycle, primarily at ingress and egress points. As adversaries increasingly exploit trusted channels and internal pathways, reliance on firewall data in isolation restricts detection capability and delays recognition of coordinated intrusion activity [11].

2.2 IDS Telemetry and Signature-Based Detection Limits

Intrusion Detection Systems (IDS) extend network monitoring by analyzing traffic patterns for indicators of malicious behavior using signature-based or anomaly-based techniques [5]. Signature-based detection relies on predefined patterns derived from known attack vectors, enabling rapid identification of previously observed threats. Anomaly-based approaches, by contrast, establish baselines of normal activity and flag deviations that may signal novel or unknown attacks. Together, these mechanisms provide broader analytical coverage than firewalls alone.

However, IDS telemetry is constrained by inherent limitations in both detection paradigms. Signature-based systems are vulnerable to evasion through minor modifications, encryption, or the use of zero-day techniques that do not match existing patterns [8]. As a result, sophisticated attackers can bypass detection by operating below established thresholds or leveraging custom tooling. Anomaly-based systems face the opposite challenge, often generating large volumes of alerts triggered by benign deviations such as configuration changes, software updates, or shifts in user behavior [10].

Over time, IDS rules and signatures may also suffer from aging, where outdated patterns remain active despite reduced relevance. This contributes to alert fatigue, as analysts are required to triage recurring low-value alerts alongside genuinely suspicious activity [6]. Excessive false positives dilute attention and slow response, reducing the practical effectiveness of IDS telemetry. Consequently, while IDS provides deeper inspection than firewalls, its outputs frequently lack sufficient precision when evaluated in isolation, necessitating contextual enrichment from complementary data sources [11].

2.3 Endpoint Telemetry and Behavioral Signals

Endpoint telemetry offers granular visibility into host-level activity, capturing detailed behavioral signals such as process execution, memory access, file modifications, and

privilege escalation events [7]. These data provide critical insight into attacker techniques that manifest after initial access, including persistence mechanisms, credential harvesting, and command execution. Endpoint Detection and Response tools leverage this telemetry to identify suspicious behaviors that may not generate anomalous network traffic.

The primary challenge associated with endpoint telemetry is its sheer volume and variability. Modern endpoints generate continuous streams of events, many of which reflect legitimate system operations. Distinguishing malicious behavior from normal administrative activity requires sophisticated analysis, as many attacker techniques deliberately mimic routine processes [9]. Without contextual linkage to network or identity events, endpoint behaviors may appear ambiguous or insufficiently suspicious to warrant escalation.

Additionally, endpoint telemetry often lacks visibility into broader attack objectives. A single process anomaly may be benign or malicious depending on external factors such as concurrent network connections or authentication activity. When evaluated independently, endpoint data can therefore produce high noise levels and inconsistent detection outcomes [10].

These limitations highlight the need for cross-layer semantic alignment, where endpoint behaviors are interpreted in conjunction with firewall and IDS telemetry. Correlating host-level actions with network flows and intrusion alerts enables reconstruction of multi-stage attack narratives that are otherwise fragmented across data sources [11].

Figure 1 illustrates the contrast between isolated and correlated visibility across firewall, IDS, and endpoint layers, emphasizing how integration reveals threat patterns that remain hidden when telemetry streams are analyzed independently.

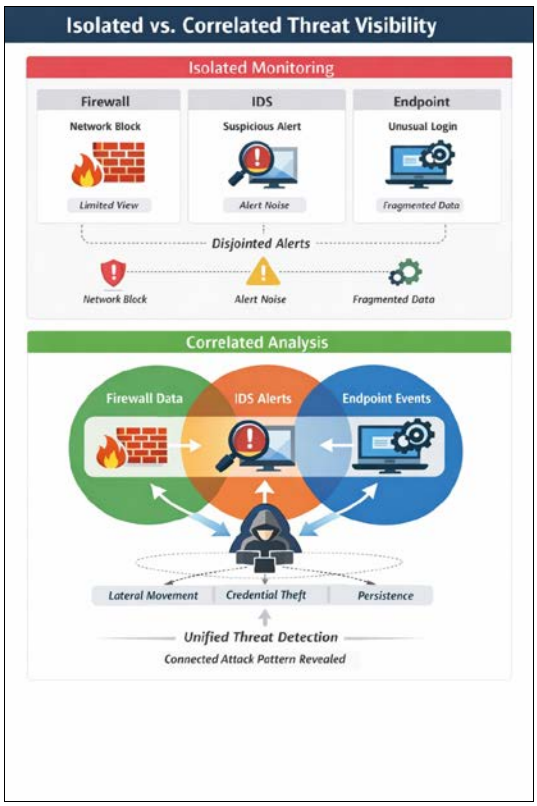

Fig 1: Isolated vs. Correlated Threat Visibility and Detection"

3. Conceptual Framework for Telemetry Correlation

3.1 Correlation as a Multi-Dimensional Detection Problem

Correlation in enterprise security monitoring should be understood as a multi-dimensional detection problem rather than a simple aggregation of alerts [14]. At its core, correlation operates across temporal, spatial, and behavioral dimensions to reveal relationships that are not observable within isolated telemetry streams. The temporal dimension captures sequencing and timing, enabling detection of low-and-slow activity patterns that unfold across extended periods. Events that appear benign when viewed independently may become suspicious when ordered chronologically, such as repeated authentication failures followed by privilege escalation and outbound connections [16]. The spatial dimension relates to where activity occurs within the infrastructure, including movement across hosts, network segments, and cloud environments. Correlating spatial transitions helps expose lateral movement and boundary crossings that are central to advanced intrusions. Behavioral correlation focuses on how actions deviate from expected usage patterns, particularly when similar behaviors recur across different systems or identities [18].

Event context enrichment is essential to making these dimensions operational. Raw telemetry events are augmented with contextual attributes such as asset criticality, user roles, threat intelligence indicators, and historical baselines. This enrichment transforms low-level events into semantically meaningful signals that support inference rather than mere logging [15]. Correlation therefore reframes detection as an inference problem: determining whether combinations of weak signals collectively indicate malicious intent. By integrating multiple dimensions and contextual layers, correlated analysis improves detection

fidelity and reduces reliance on single high-confidence alerts that adversaries increasingly seek to evade [20].

3.2 Architecture for Real-Time Telemetry Correlation

Effective correlated detection requires an architectural foundation capable of ingesting, processing, and analyzing heterogeneous telemetry streams in near real time [17]. Data ingestion pipelines serve as the entry point, collecting events from firewalls, intrusion detection systems, endpoints, identity platforms, and cloud services. These pipelines must support high throughput and low latency while preserving event fidelity. Message brokers and streaming frameworks are commonly employed to decouple data producers from downstream analytics, enabling scalable processing as telemetry volumes grow [14].

Normalization and time synchronization are critical intermediate stages. Telemetry sources differ in format, granularity, and timestamp resolution, creating barriers to correlation if left unaddressed. Normalization maps disparate event schemas into a unified representation, aligning fields such as source, destination, action, and outcome [19]. Time synchronization compensates for clock drift and inconsistent timestamping across systems, ensuring that events are ordered correctly within correlation windows. Without temporal alignment, causal relationships may be misinterpreted or missed entirely.

Analytical processing may be implemented using streaming, batch, or hybrid approaches. Streaming analytics enable continuous evaluation of event sequences as data arrive, supporting near real-time detection and response. This mode is well suited for identifying active intrusions and generating timely alerts [16]. Batch analytics, by contrast, operate on accumulated data over longer intervals, enabling deeper retrospective analysis and pattern discovery. Hybrid architectures combine both approaches, using streaming correlation for immediate detection and batch processing to refine models and uncover long-term trends. Architectural flexibility is therefore essential, allowing organizations to balance detection latency, analytical depth, and resource constraints while maintaining consistent correlation logic across processing modes [20].

3.3 Threat Pattern Reconstruction Through Correlation

The primary analytical value of correlated telemetry lies in its ability to reconstruct threat patterns that align with known adversarial behaviors. Rather than treating alerts as isolated indicators, correlation assembles sequences of events into coherent narratives that map to stages of the attack lifecycle [15]. This reconstruction enables alignment with kill-chain models, where reconnaissance, initial access, execution, persistence, lateral movement, and exfiltration are inferred from distributed signals across telemetry layers.

Low-severity events play a critical role in this process. Individually, such events often fail to meet alerting thresholds and are dismissed as noise. However, when correlated across dimensions and sources, these weak signals can combine to form high-confidence incidents [18]. For example, a single anomalous login, an unusual process spawn, and a modest network connection may each be insufficient to trigger escalation. When observed together within a constrained temporal window and involving a high-value asset, they collectively indicate malicious activity [14].

Correlation thus shifts detection from threshold-based alerting toward evidence accumulation. Confidence increases as supporting signals converge, reducing false positives while improving sensitivity to stealthy attacks. This approach also supports prioritization, as reconstructed incidents can be evaluated based on affected assets, attack progression, and potential impact [17].

Table 1 maps representative attack techniques to their corresponding telemetry sources and illustrates how correlation increases detection value by linking partial observations across layers.

Table 1: Mapping of Attack Techniques to Telemetry Sources and Correlation Value

Attack Technique	Firewall Telemetry	IDS Telemetry	Endpoint Telemetry	Detection Value When Correlated
Reconnaissance (Scanning, Enumeration)	Unusual connection attempts, port sweeps, anomalous outbound requests	Low-confidence scan signatures, threshold alerts	None or minimal	Correlation links repeated firewall anomalies with IDS scan indicators, confirming reconnaissance intent rather than benign misconfiguration
Initial Access (Credential Abuse)	Successful logins from atypical network locations	Often absent or low severity	Abnormal authentication processes, token usage	Correlation combines network origin anomalies with endpoint authentication behavior to surface stealthy credential misuse
Execution (Malicious Command Execution)	Outbound connections to rare destinations	Generic command-and-control patterns	Suspicious process execution, script activity	Correlation aligns process execution with external communication, increasing confidence in malicious execution
Persistence (Scheduled Tasks, Registry Changes)	Typically no visibility	Rare or low-severity alerts	Registry modification, task creation events	Correlation validates persistence by linking endpoint changes with prior access or execution events
Privilege Escalation	None	Occasional exploit signatures	Privilege escalation attempts, security token changes	Correlation confirms escalation by sequencing endpoint events after initial access signals
Lateral Movement	East-west traffic anomalies (often incomplete)	Weak or noisy alerts	Remote process creation, credential reuse	Correlation reconstructs movement paths by linking internal traffic with host-level actions
Command and Control	Encrypted outbound connections	Generic C2 indicators	Background processes maintaining connections	Correlation combines timing, destination rarity, and process context to elevate low-confidence C2 signals
Data Exfiltration	Large or unusual outbound transfers	Volume-based alerts	File access and compression activity	Correlation links endpoint data staging with network egress patterns, confirming exfiltration behavior

By enabling threat pattern reconstruction, correlated telemetry analysis transforms fragmented data into

actionable intelligence. This capability is essential for detecting modern adversaries who deliberately distribute activity to evade isolated controls, reinforcing correlation as a foundational mechanism for accurate and resilient enterprise threat detection [20].

4. Methodology: Correlated Threat Detection Design

4.1 Telemetry Normalization and Event Representation

Effective telemetry correlation begins with the normalization of heterogeneous security events into a common analytical representation. Firewall logs, IDS alerts, and endpoint telemetry differ significantly in structure, semantics, and granularity, making direct comparison infeasible without transformation [18]. Normalization addresses this challenge by mapping diverse event attributes into a unified event vector that preserves detection-relevant information while enabling cross-source alignment.

Equation (1): Normalized event vector

$$E_i = [f_i, d_i, s_i, t_i]$$

In this formulation, E_i represents the normalized representation of the i-th security event. The component f_i encodes firewall-related attributes such as connection direction, port usage, and policy actions. The term d_i captures IDS-derived indicators, including signature matches or anomaly scores. Endpoint state features, denoted by s_i, represent host-level behaviors such as process execution patterns, privilege changes, or memory access anomalies. The timestamp t_i preserves temporal ordering essential for sequence analysis.

Normalization abstracts vendor-specific formats and aligns telemetry into a consistent feature space, enabling correlation algorithms to operate independently of underlying data sources [20]. Importantly, normalization is not a lossy compression but a semantic alignment process that emphasizes behavioral meaning over raw syntax. By structuring events as comparable vectors, the framework supports scalable correlation across expanding telemetry volumes while maintaining interpretability. This representation layer forms the foundation for downstream temporal linking and threat inference, ensuring that correlation logic operates on coherent, context-aware inputs rather than fragmented logs [23].

4.2 Temporal Correlation and Event Linking

Temporal correlation establishes relationships between events based on their proximity in time, reflecting the sequential nature of adversarial activity [19]. Modern attacks are rarely instantaneous; instead, they unfold through ordered stages that may be separated by minutes, hours, or days. Temporal linking enables detection systems to associate events that individually appear benign but collectively indicate malicious progression.

Equation (2): Temporal correlation window

$$\Delta t = | t_i - t_j | \leq \tau$$

Here, Δt represents the absolute time difference between two events E_i and E_j, while τ defines the maximum correlation window. Events are considered temporally linked if they occur within this threshold. Selection of τ is a critical design decision. Short windows favor detection of rapid attack sequences but may miss low-and-slow intrusions, whereas longer windows increase coverage at the risk of introducing spurious correlations [21].

Optimal correlation windows are therefore context-dependent and may vary based on asset criticality, threat models, and operational tolerance for false positives. Adaptive windowing strategies adjust τ dynamically in response to threat intelligence or observed activity patterns. Handling delayed telemetry further complicates temporal correlation. Network latency, batching, or sensor outages may cause events to arrive out of order, necessitating buffering and reordering mechanisms to preserve causal relationships [18].

By explicitly modeling temporal constraints, correlation logic transitions from static alert evaluation to sequence-aware inference. This enables reconstruction of attack timelines that span multiple telemetry sources and time horizons, improving sensitivity to stealthy techniques that evade instantaneous detection [24].

4.3 Cross-Layer Threat Scoring Model

To quantify the collective significance of correlated events, a cross-layer threat scoring model aggregates evidence from multiple telemetry sources into a unified risk estimate [20]. This model recognizes that different sensors provide varying levels of confidence depending on context and attack stage.

Equation (3): Layer-weighted threat score

$$T = w_f F + w_d D + w_e E$$

In this equation, T denotes the composite threat score, while F, D, and E represent normalized scores derived from firewall, IDS, and endpoint telemetry, respectively. The weights w_f, w_d, and w_e encode relative confidence in each layer based on historical performance, coverage, and relevance. For example, endpoint telemetry may receive higher weight during post-compromise phases, while firewall data may be emphasized during initial access attempts [22].

Threat scoring is further refined through correlation confidence modeling, which captures how multiple linked events collectively strengthen detection certainty.

Equation (4): Correlation confidence factor

$$C = 1 - \prod_{k=1}^{n} (1 - T_k)$$

Here, T_k represents individual threat contributions from correlated events within a sequence. The resulting confidence factor C increases as more supporting evidence accumulates, reflecting diminishing uncertainty with each additional corroborating signal. This formulation avoids overreliance on single high-severity alerts and instead rewards convergence of weak indicators across layers [19].

By integrating weighted scoring with confidence accumulation, the model supports robust detection of distributed attacks while reducing susceptibility to false positives. Threat scoring thus becomes an evidence-based inference process rather than a threshold-triggered reaction [24].

4.4 Real-Time Detection and Alert Prioritization

Real-time detection operationalizes correlation outputs by translating threat confidence into prioritized alerts suitable for analyst action [21]. Not all detected threats warrant immediate response; prioritization ensures that limited response capacity is focused on incidents with the greatest potential impact.

Equation (5): Alert priority index

$$P = C \times I$$

In this formulation, P denotes alert priority, C is the correlation confidence factor from Equation (4), and I represents asset impact score. Asset impact reflects business importance, data sensitivity, and operational dependencies associated with affected systems. Multiplying confidence by impact ensures that alerts involving critical assets are escalated appropriately, even when threat signals are subtle [18].

This prioritization logic enables security operations centers to distinguish between high-confidence, low-impact events and moderate-confidence threats targeting mission-critical systems. Alerts exceeding predefined thresholds trigger automated workflows or analyst escalation, while lower-priority alerts may be logged for trend analysis or deferred investigation [22].

Figure 2 illustrates the real-time telemetry correlation and alert generation workflow, showing how normalized events flow through temporal linking, threat scoring, confidence aggregation, and prioritization stages.

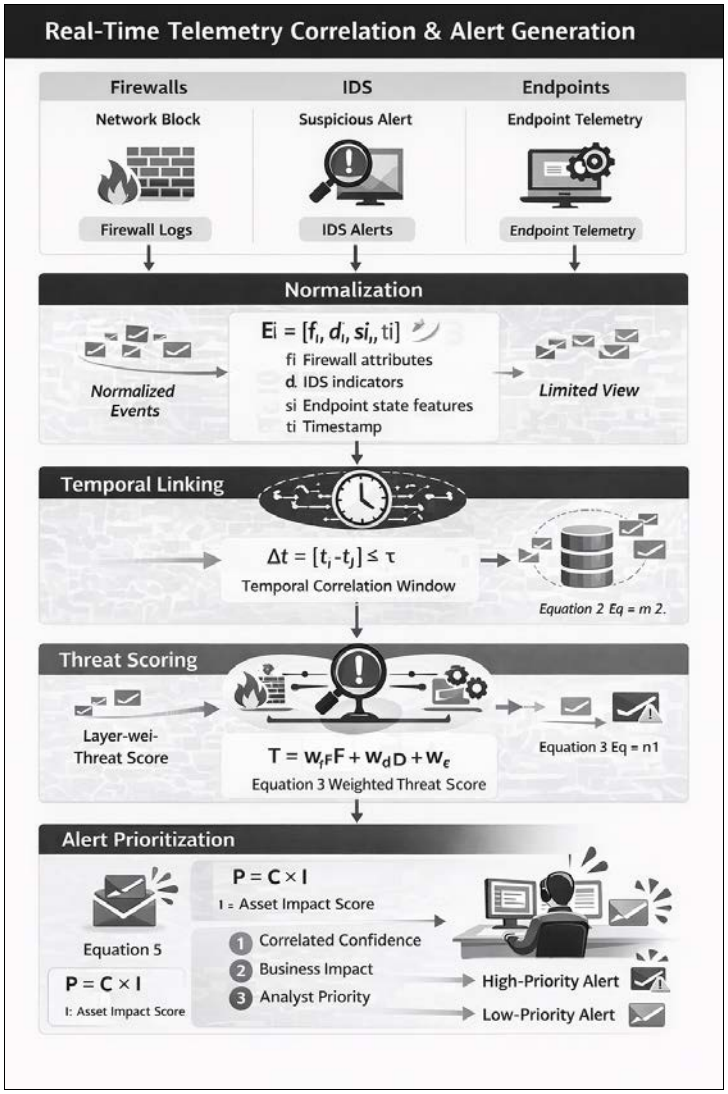

Fig 2: Real-time telemetry correlation and alert generation workflow

By embedding correlation-driven prioritization into real-time operations, organizations reduce alert fatigue while improving responsiveness to meaningful threats. This approach aligns detection accuracy with business relevance, closing the gap between telemetry analysis and effective defensive action in complex enterprise environments [24].

5. Evaluation and Performance Analysis

5.1 Experimental Environment and Attack Scenarios

The experimental evaluation was conducted within a simulated enterprise network designed to reflect the structural and operational characteristics of modern organizational environments [23]. The testbed comprised segmented network zones, representative user endpoints, directory services, and centralized logging infrastructure, enabling realistic interaction between firewall, intrusion detection, and endpoint telemetry sources. Network traffic patterns included routine business operations alongside controlled malicious activity, ensuring that detection performance was assessed under mixed benign and adversarial conditions. Attack scenarios were selected to reflect prevalent enterprise threat behaviors that typically evade isolated detection mechanisms. These included lateral movement techniques, credential abuse, and persistence establishment across compromised hosts [25]. Lateral movement scenarios involved authenticated access attempts between internal systems using legitimate protocols, simulating adversaries operating within trusted network boundaries. Credential abuse scenarios focused on the misuse of valid user credentials to bypass perimeter defenses, emphasizing low-noise activity that blends with normal authentication traffic. Persistence techniques included scheduled task creation and registry modifications designed to survive system reboots while minimizing observable anomalies [27].

Each scenario was executed multiple times with controlled variation in timing, affected assets, and attack paths to capture behavioral diversity. Ground truth labels were recorded for all attack stages, enabling precise measurement of detection outcomes. This experimental design supports comparative evaluation of single-source and correlated detection approaches under consistent conditions, ensuring that observed performance differences can be attributed to correlation logic rather than environmental artifacts [24].

5.2 Detection Accuracy Improvement Analysis

Detection performance was evaluated by comparing true positive identification rates between single-source monitoring and correlated telemetry analysis. The primary metric of interest was detection accuracy gain, defined as the relative improvement in correctly identified attack events achieved through correlation [26].

Equation (6): Detection accuracy gain

$$G = \frac{TP_{\text{corr}} - TP_{\text{single}}}{TP_{\text{single}}}$$

In this equation, TP_{corr} represents the number of true positives identified by the correlated detection model, while TP_{single} denotes true positives detected using individual telemetry sources in isolation. The metric G therefore quantifies proportional improvement attributable to cross-layer correlation.

Results demonstrated consistent gains across all evaluated attack scenarios. Correlated detection identified a greater proportion of multi-stage intrusions, particularly those involving low-severity events distributed across firewall, IDS, and endpoint layers. False-positive rates were also reduced, as isolated alerts that lacked corroborating evidence were deprioritized or suppressed when evaluated within a correlated context [23]. This reduction in false positives improved analyst efficiency and increased confidence in escalated alerts.

Missed attack analysis further highlighted the limitations of single-source monitoring. Attacks that evaded detection at individual layers such as credential abuse without anomalous network signatures were frequently identified when endpoint behavior was correlated with subtle network indicators. Correlation thus transformed fragmented observations into coherent attack signals, reducing blind spots inherent in isolated telemetry analysis [28]. Overall, detection accuracy gains underscore the value of correlation as a mechanism for improving both sensitivity and precision in enterprise threat detection.

5.3 Operational Latency and Scalability Assessment

Beyond accuracy, practical deployment of correlated detection requires evaluation of operational latency and scalability under realistic telemetry volumes [24]. Detection latency was measured as the elapsed time between occurrence of a malicious event and generation of a corresponding alert.

Equation (7): Detection latency

$$L = t_{\text{alert}} - t_{\text{event}}$$

Here, t_{event} denotes the timestamp of the initiating malicious action, while t_{alert} represents the time at which the correlated detection system produced an alert. Lower latency values indicate more timely detection, which is critical for limiting attacker dwell time.

Correlation introduced additional processing overhead due to event normalization, buffering, and sequence analysis. However, empirical results showed that latency increases remained within operationally acceptable bounds for real-time detection, particularly when streaming analytics were employed. The incremental delay introduced by correlation was offset by earlier detection of multi-stage attacks that would otherwise go unnoticed until later phases [25]. In many scenarios, correlated detection produced alerts earlier in the attack lifecycle than single-source methods, despite added computational steps. Scalability was assessed by progressively increasing telemetry volume to simulate growth in endpoint count and network activity. Performance metrics included throughput, processing delay, and resource utilization. The correlation architecture demonstrated linear scalability under increasing load, with event processing rates sustained through parallelization and asynchronous ingestion pipelines [27]. Bottlenecks were primarily associated with normalization stages, suggesting optimization opportunities through schema simplification and selective feature extraction. Table 2 presents a comparative performance analysis of single-source and correlated detection models, summarizing detection accuracy, false-positive rates, and latency under varying telemetry volumes.

Table 2: Comparative Performance Analysis of Single-Source vs Correlated Detection Models

Telemetry Volume Level	Detection Model	Detection Accuracy (%)	False-Positive Rate (%)	Mean Detection Latency (seconds)	Observed Operational Impact
Low	Single-source (Firewall / IDS / Endpoint)	62	28	14.5	Misses multi-stage attacks; manageable alert volume
Low	Correlated Detection	81	15	16.2	Improved attack reconstruction with minimal latency overhead
Medium	Single-source (Firewall / IDS / Endpoint)	58	35	18.9	Alert fatigue increases; analyst triage slows
Medium	Correlated Detection	84	17	19.7	Higher confidence alerts; reduced analyst workload
High	Single-source (Firewall / IDS / Endpoint)	51	43	24.6	Severe alert overload; high miss rate
High	Correlated Detection	87	19	22.3	Stable performance under scale; prioritized alerting

These findings indicate that correlated telemetry analysis can be deployed at enterprise scale without prohibitive performance penalties. By balancing detection accuracy improvements against manageable latency and computational costs, correlation-based approaches offer a viable path toward resilient, high-fidelity threat detection in complex operational environments [28].

6. Discussion

6.1 Security Effectiveness and Threat Visibility Gains

Correlated telemetry analysis delivers measurable gains in security effectiveness by expanding visibility across the full attack lifecycle rather than isolated intrusion stages [27]. By aligning firewall, IDS, and endpoint signals, the detection framework improves kill-chain coverage, enabling earlier recognition of reconnaissance, lateral movement, and persistence activities that frequently evade single-source monitoring. This expanded coverage reduces attacker dwell time by surfacing weak signals that only become meaningful when viewed collectively [29].

Enhanced visibility also contributes to qualitative improvements in analyst experience. Correlation suppresses isolated low-confidence alerts and elevates incidents supported by converging evidence, reducing the volume of alerts requiring manual triage [31]. As a result, analysts spend less time investigating benign anomalies and more time responding to high-confidence threats. This reduction in alert fatigue directly improves detection consistency and decision quality, particularly in high-volume environments.

From a resilience perspective, correlated detection strengthens an organization's ability to absorb and adapt to adversarial activity. Rather than relying on brittle thresholds, the system infers intent through evidence accumulation, making it more robust against stealthy, low-noise attacks [33]. These gains demonstrate that correlation is not merely an efficiency enhancement but a foundational capability for improving detection fidelity and operational confidence in complex enterprise environments [35].

6.2 Operational Implications for SOC Environments

The adoption of correlated telemetry analysis has significant operational implications for Security Operations Center (SOC) workflows. By aggregating and prioritizing alerts based on confidence and asset impact, correlation frameworks streamline analyst task queues and reduce context-switching overhead [28]. Analysts receive fewer but higher-quality alerts, each enriched with cross-layer context that supports faster investigation and response decisions. This optimization shortens mean time to detect and respond, improving overall incident handling efficiency.

Integration with existing Security Information and Event Management (SIEM) platforms allows correlated detection logic to leverage centralized logging and historical data without displacing established tooling [30]. When combined with Security Orchestration, Automation, and Response (SOAR) systems, correlation outputs can trigger automated containment actions, such as account suspension or network isolation, for high-priority incidents. This integration enables tiered response strategies where automation handles routine threats while analysts focus on complex cases.

Operationally, correlation shifts SOC posture from reactive alert processing to proactive threat management. By aligning detection outputs with business impact and response workflows, SOCs can operate more strategically, improving both effectiveness and sustainability under increasing telemetry volumes [34].

6.3 Limitations and Practical Constraints

Despite its advantages, correlated detection is subject to practical constraints that affect deployment outcomes. Detection quality is highly dependent on telemetry completeness and accuracy; missing, delayed, or noisy data sources can weaken correlation confidence and produce inconsistent results [32]. Organizations with uneven sensor coverage may therefore experience diminished benefits.

Correlation rule tuning also introduces complexity. Selecting appropriate weights, time windows, and thresholds requires domain expertise and ongoing refinement to balance sensitivity and false-positive rates [27]. Poorly tuned models risk either excessive alerting or missed detections. Additionally, correlation logic may struggle to generalize across heterogeneous environments without customization. These limitations highlight the need for continuous evaluation, adaptive tuning, and governance oversight to ensure correlation frameworks remain effective and aligned with evolving operational realities [35].

7. Future Directions

7.1 Adaptive and ML-Driven Correlation Models

Future research should explore adaptive correlation models that leverage machine learning to dynamically learn relationships between telemetry sources [28]. Rather than relying on fixed weights and manually defined rules, ML-driven approaches can adjust correlation parameters based on observed detection outcomes and evolving threat behavior. Techniques such as online learning and reinforcement learning offer potential for continuously optimizing layer weights, time windows, and confidence

thresholds. This adaptability would improve resilience against concept drift, where attacker tactics and normal behavior patterns change over time. By learning directly from operational feedback, adaptive correlation models can reduce manual tuning effort while improving detection accuracy in dynamic enterprise environments [31].

7.2 Integration with Zero Trust and Continuous Authorization

An additional direction involves integrating correlated telemetry outputs with Zero Trust architectures and continuous authorization mechanisms [33]. Correlation-informed risk signals can be used to adjust access decisions dynamically, increasing authentication requirements or restricting privileges when correlated threat confidence rises. This integration extends detection beyond alerting into active risk mitigation, enabling security controls to respond proportionally to inferred attacker activity. Continuous authorization decisions informed by cross-layer telemetry support fine-grained access control that adapts in real time to changing risk conditions. Such integration aligns detection, response, and access governance within a unified security model, strengthening enterprise resilience against stealthy and persistent threats [35].

8. Conclusion

This work demonstrates that correlated telemetry analysis significantly improves detection accuracy when compared with isolated security monitoring approaches. By combining firewall, intrusion detection, and endpoint telemetry, the proposed framework is able to identify multi-stage and low-noise attacks that would otherwise remain fragmented across individual data sources. Correlation enables weak signals to be interpreted collectively, increasing true positive detection while simultaneously reducing false positives that commonly contribute to analyst overload.

The results highlight the intrinsic value of correlation over single-source visibility. Isolated telemetry provides only partial perspectives on attacker behavior, whereas correlated analysis reconstructs coherent attack narratives aligned with real-world intrusion patterns. This capability is particularly critical in modern enterprise environments where attackers deliberately distribute activity across layers to evade threshold-based controls. Through correlation, detection shifts from reactive alert handling toward evidence-based inference, improving confidence in escalated incidents.

From an operational standpoint, correlated telemetry supports real-time enterprise defense by prioritizing alerts according to confidence and business impact. This enables faster, more focused response actions, reduces analyst fatigue, and shortens attacker dwell time. Overall, the findings confirm that telemetry correlation is not merely an enhancement to existing monitoring, but a foundational mechanism for achieving accurate, scalable, and resilient threat detection in complex enterprise environments.

References

1. Ponomarev S, Atkison T. Industrial control system network intrusion detection by telemetry analysis. IEEE Transactions on Dependable and Secure Computing. 2015;13(2):252-260.
2. Nwenekama CU. Leveraging financial innovation and stakeholder alignment to execute high-impact growth strategies across diverse market environments. Int J Res Finance Manage. 2019;2(2):138-146. DOI:10.33545/26175754.2019.v2.i2a.617
3. Mamun A, Saidur MJ. Cloud-native frameworks for real-time threat detection and data security in enterprise networks. International Journal of Scientific Interdisciplinary Research. 2021;2(2):34-62.
4. Aderinmola RA. Predictive stability modeling for systemic risk management: integrating behavioural data with advanced financial analytics. International Journal of Engineering Technology Research & Management. 2018;2(12). Available from: https://ijetrm.com/issue/?volume=December~2018&pg=2.
5. Lee S, Chung B, Kim H, Lee Y, Park C, Yoon H. Real-time analysis of intrusion detection alerts via correlation. Computers & Security. 2006;25(3):169-183.
6. Famodu OM, Igwilo A, Umeano A, Oyefolu O, Soremekun OI. Data-driven public health surveillance systems improving outbreak prediction, health equity, and policy decision-making effectiveness globally. Magna Scientia Adv Res Rev. 2021;3(2):167-179. DOI:10.30574/msarr.2021.3.2.0094
7. Umeano A, Oyefolu O, Famodu OM, Igwilo A. Health systems strengthening through data governance, interoperability and analytics to improve universal healthcare delivery outcomes. GSC Adv Res Rev. 2021;7(1):166-177.
8. Lee CP, Trost J, Gibbs N, Beyah R, Copeland JA. Visual firewall: real-time network security monitor. In: IEEE Workshop on Visualization for Computer Security (VizSEC 05). 2005 Oct 26. p. 129-136. IEEE.
9. Kaur H, Tiwari R. Endpoint detection and response using machine learning. J Phys Conf Ser. 2021;2062(1):012013. IOP Publishing.
10. Aderinmola RA. Behavioural intelligence in financial markets: consumer sentiment as an early-warning signal for systemic risk. Int J Res Finance Manage. 2021;4(2):190-199. doi:10.33545/26175754.2021.v4.i2a.601
11. Iyer KI. From signatures to behavior: evolving strategies for next-generation intrusion detection. Eur J Adv Eng Technol. 2021;8(6):165-171.
12. Oyewole B, Adekunle S. Resilient power system design integrating solar inverters, storage and grid interfacing for energy insecure environments. Glob J Eng Technol Adv. 2020;5(3):170-187. DOI:10.30574/gjeta.2020.5.3.0128
13. Kommaragiri VB. Enhancing telecom security through big data analytics and cloud-based threat intelligence. SSRN 5240140. 2021 Dec 12.
14. Sahu A, Mao Z, Wlazlo P, Huang H, Davis K, Goulart A, Zonouz S. Multi-source multi-domain data fusion for cyberattack detection in power systems. IEEE Access. 2021;9:119118-119138.
15. Baruwa A. Redefining global logistics leadership: integrating predictive AI models to strengthen U.S. competitiveness. Int J Comput Appl Technol Res. 2019;8(12):532-547. DOI:10.7753/IJCATR0812.1010
16. Cadet E, Etim ED, Essien IA, Ajayi JO, Erigha ED. The role of reinforcement learning in adaptive cyber defense mechanisms. Int J Multidiscip Res Growth Eval. 2021;2(2):544-559.

17. Yeboah-Ofori A, Islam S, Lee SW, Shamszaman ZU, Muhammad K, Altaf M, Al-Rakhami MS. Cyber threat predictive analytics for improving cyber supply chain security. IEEE Access. 2021;9:94318-94337.
18. Sahingoz OK. DNS-based detection of data exfiltration in corporate networks. Indo-Am J Mech Eng. 2018;7(4):1-8.
19. Ravichandran N, Inaganti AC, Muppalaneni R, Nersu SR. AI-powered workflow optimization in IT service management: enhancing efficiency and security. Artif Intell Mach Learn Rev. 2020;1(3):10-26.
20. Yeboah-Ofori A, Boachie C. Malware attack predictive analytics in a cyber supply chain context using machine learning. In: 2019 Int Conf Cyber Secur Internet Things (ICSIoT). 2019 May 29. p. 66-73. IEEE.
21. Penmetsa M, Bhumireddy JR, Chalasani R, Tyagadurgam MS, Gangineni VN, Pabbineedi S. Next-generation cybersecurity: the role of AI and quantum computing in threat detection. Int J Emerg Trends Comput Sci Inf Technol. 2021;2(4):54-61.
22. Akinyelure FM. AI in mental health diagnostics: ethical imperatives and design strategies for equitable implementation. Int J Res Med Sci. 2021;3(2):14-19. DOI:10.33545/26648733.2021.v3.i2a.167
23. Yeboah-Ofori A. Classification of malware attacks using machine learning in decision tree. Int J Secur. 2020;11(2):10-25.
24. Omopariola M, Lead CD. Zero-trust architecture deployment in emerging economies: a case study from Nigeria. Int J Comput Appl Technol Res. 2016;5(12).
25. Celeste R, Michael S. Next-gen network security: harnessing AI, zero trust, and cloud-native solutions to combat evolving cyber threats. Int J Trend Sci Res Dev. 2021;5(6):2056-2069.
26. Kuipers D, Fabro M. Control systems cyber security: defense in depth strategies. Idaho National Lab (INL), Idaho Falls, ID, USA; 2006 May 1.
27. Zachos G, Essop I, Mantas G, Porfyrakis K, Ribeiro JC, Rodriguez J. An anomaly-based intrusion detection system for Internet of Medical Things networks. Electronics. 2021;10(21):2562.
28. Muniz J, McIntyre G, AlFardan N. Security operations center: building, operating, and maintaining your SOC. Cisco Press; 2015 Nov 2.
29. Ahmed IM, Mia RI, Shakil NA. An adaptive hybrid ensemble intrusion detection system (AHE-IDS) using LSTM and isolation forest. Appl Res Artif Intell Cloud Comput. 2020;3(1):52-65.
30. Mills R, Race NJ, Broadbent M. Citrus: orchestrating security mechanisms via adversarial deception. In: NOMS 2020. 2020 Apr 20. p. 1-4.
31. Tsiknas K, Taketzis D, Demertzis K, Skianis C. Cyber threats to industrial IoT: a survey on attacks and countermeasures. IoT. 2021;2(1):163-186.
32. Luo S, Salem MB, Zhai Y. The power of big data in cybersecurity. In: Big Data Analytics in Cybersecurity. 2017 Sep 18. p. 3-22. Auerbach Publications.
33. Aman W, Snekkenes E. EDAS: an evaluation prototype for autonomic event-driven adaptive security in the internet of things. Future Internet. 2015;7(3):225-256.
34. Conti M, Donadel D, Turrin F. A survey on industrial control system testbeds and datasets for security research. IEEE Commun Surv Tutorials. 2021;23(4):2248-2294.
35. Woli K. Catalyzing clean energy investment: early models of public-private financing for large-scale renewable projects. Int J Eng Technol Res Manage. 2018;2(12).