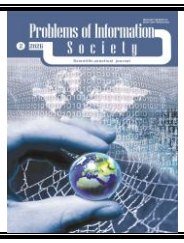

www.jpis.az

Real time attack visualization map for security operation centers

Rashad Aliyev[1], *Sarfi Habibova*[2]

Khazar University, Baku, Azerbaijan

[1]rashad@aliev.info, [2]serfihebibova@gmail.com

[1]orcid.org/0009-0000-4554-0084, [2]orcid.org/0009-0008-5691-3585

ARTICLE INFO

https://doi.org/10.25045/jpis.v17.i2.12

Article history:
Received 16 February 2026
Received in revised form
04 May 2026
Accepted 30 June 2026

Keywords:
Cyber security
Threat maps
Security alerts
Cyber-attack visualization
Security Operation Centers
Intrusion Detection Systems

ABSTRACT

Security Operations Centers commonly visualize alerts as text logs or static dashboards, which impairs situational awareness in the case of high-volume attacks. This paper describes a real-time visualization layer built on top of an on-premises Wazuh and Suricata deployment that converts incoming security alerts into animated geographic arcs on an interactive map, using the source IP of each event to estimate its country of origin. The system is composed entirely of freely available components and operates within the organization's own network, so raw event data never leaves local infrastructure. The proposed system differs from public threat maps provided by commercial vendors, which are based on aggregated and anonymized feeds that do not correspond to any particular network. Instead, the system provides a shared and immediately readable picture of ongoing activity to analysts and non-technical stakeholders, displaying only alerts generated by the monitored infrastructure.

1. Introduction

Security Information and Event Management (SIEM) platforms have become a standard component of Security Operations Centers (SOC) operations worldwide, and open-source options such as Wazuh are increasingly used to collect and correlate events from servers, applications, and network devices into a single repository. Network-level tools such as Suricata complement this by inspecting live traffic and forwarding detections into the same pipeline, giving analysts broader coverage. The scale of the underlying problem is considerable: combined vulnerability databases like CVE, NVD, BID, and SYM list well over half a million known entries, showing how broad the surface is that SOC teams are expected to monitor (Aliyev & Peñalver, 2016). Attack surfaces of this size are not limited to conventional IT networks; surveys of cyber-physical systems report a comparable growth in attack diversity as operational technology and IT infrastructure converge, reinforcing the need for defenses that scale with data volume rather than analyst headcount (Duo et al., 2022).

Despite this, most SIEM interfaces still present events as scrolling text, tables, and fixed charts. Under heavy alert volume, a condition commonly referred to as alert fatigue, this format makes it difficult for analysts to reconstruct attack sequences in a timely fashion, delaying decisions and lengthening response times. Prior work has argued that visual presentation of security events reduces cognitive load and accelerates pattern recognition relative to text-first interfaces (Bhatt et al., 2014), a conclusion consistent with research on the human factor in cybersecurity, which identifies analyst attention and workload as a determining variable in how effectively technical controls actually perform in practice (Mahmudova, 2024). SIEM, together with intrusion detection/prevention systems and log-management tooling, remains one of the foundational layers of a Blue Team's toolkit,

which is exactly why the interface built around it matters (Aliyev, 2025).

A separate category of tools, commercial live cyber-threat maps such as the Kaspersky Cyber Threat Map, the Radware Live Threat Map, and Check Point's ThreatMap, already displays attacks geographically, and open community projects such as Global Threat Map or geoip-attack-map visualizers demonstrate the same idea using public or honeypot-sourced feeds. These platforms are compelling to look at, but because their data is aggregated and anonymized across many unrelated networks, none of them reflect events actually targeting a specific organization, and none can be used for operational decisions about that organization's own infrastructure.

This creates a gap between two otherwise mature categories of tooling: SIEM systems that hold detailed, organization-specific event data but display it as text, and threat maps that display geography but not an organization's own data. Closing this gap is of particular relevance in settings such as the digital infrastructure of Azerbaijan, where in-house, open-source SOC tooling is prevalent and reliance on foreign commercial dashboards poses cost and data-sovereignty issues; recent work on AI-enabled, locally operated monitoring systems argues that keeping analytical capability inside national or organizational infrastructure is itself a component of cyber sovereignty, not only a cost consideration (Alakbarova, 2026). It also matters for communication: an interactive map of source country, IP, rule triggered, and attack type is much easier for non-technical managers to read at a glance than a table of raw log lines, and individual-level understanding of security information is itself an established gap between technical staff and decision-makers rather than a purely technical shortcoming.

2. Related work

Human-centered studies of SOC interfaces argue that effective displays should align with how analysts naturally perceive color, shape, and layout, so that attention is guided rather than demanded (Seong et al., 2020). Reviews of visual syntax used for representing attacks, including attack graphs and attack trees, similarly emphasize that the choice of visual encoding is not a cosmetic detail: different syntaxes make different aspects of an attack easy or hard to see, and a representation chosen for its analytical completeness is not automatically the one best suited to fast comprehension by a mixed audience (Lallie et al., 2020). Automated visualization pipelines built directly on top of detection systems have been shown to shorten the time between raw alert and comprehensible picture compared with manual dashboard construction, supporting the case for tightly coupling detection and visualization rather than treating them as separate stages (Alhaidari et al., 2022).

Surveys of deployed SOC visualization tooling report that combining several complementary views, geographic, temporal, and flow-based, tend to outperform any single view in isolation, though integration with standard SIEM back-ends and scalability under high traffic remain open problems. Studies focused specifically on CISO-level dashboards find that a global view of risk combined with a clear severity rating shortens the time needed to reach a decision, and that geographic context is one of the more consistently useful elements in that view. Complementary engineering-focused work has shown that Wazuh and Suricata can be integrated with lightweight external channels, such as chat-bot notifications, to shorten the time between detection and human awareness even further, confirming that the underlying detection stack is a practical foundation for real-time extensions of this kind (Jumiaty & Soewito, 2024). Nearly all this literature, however, still draws its example data from third-party or simulated feeds rather than a live, self-hosted deployment, which is the gap this work addresses directly.

It is also worth situating the choice of a geographic map against the wider set of encodings the visualization literature considers for security data. Attack graphs and attack trees, for instance, are built specifically to preserve causal and sequential relationships between steps of an intrusion, which a map cannot show; the trade-off is that reading a graph correctly requires training that a manager glancing at a screen during a demonstration simply does not have (Lallie et al., 2020). Node-link diagrams of host-to-host communication make lateral movement explicit, and treemaps compress hierarchical alert volume into a compact area, but both require the viewer to already understand the encoding convention before the display becomes informative. A geographic arc, by contrast, borrows an encoding almost everyone already knows how to read: a moving line from one place toward another. This is not evidence that maps are the superior encoding in general; automated pipelines that couple detection systems directly to graph-based

or matrix-based views have themselves been shown to shorten analyst response time in controlled settings (Alhaidari et al., 2022). Rather, it explains why a map is the right choice specifically where the interface must serve two audiences of very different technical background at once, which is the situation this system was built for.

A further point that separates this system from most of the surveyed literature is operational continuity. Much of the published work on SOC visualization evaluates prototypes against recorded, replayed, or synthetic datasets, which makes it possible to control variables cleanly but leaves open the question of how the same design performs against a live, noisy, and occasionally malformed alert stream.

This paper presents a system that renders alerts collected from an on-premises Wazuh/Suricata deployment as live geographic arcs, using only open-source components, running entirely inside the organization's own network, and requiring no external threat-intelligence feed.

The proposed system allows to identify integration issues such as inconsistent field naming between agent-generated alerts and Suricata-generated alerts, that a synthetic dataset would not expose.

3. Material and methods

3.1. Problem definition

The starting point for this work is a condition widely reported in SOC operations literature and confirmed by our own experience running an internal security operation: the number of alerts a modern SIEM generates routinely exceeds what a human analyst can read line by line. Wazuh, Suricata, firewalls, and endpoint agents each contribute their own stream, and correlation rules multiply rather than reduce the total volume once several sources report on the same underlying event. When alerts come in bursts from scanning campaigns, credential-stuffing attempts, or coordinated botnet activity, the real bottleneck is the interface the analyst is given, not the detection logic behind it.

A second, more specific problem is that the visualization options actually available to a typical on-premises SOC do not match what the research literature recommends. The tools that do offer a working geographic view, commercial threat maps from vendors such as Kaspersky, Radware, and Check Point, solve a different problem: they are built as visually persuasive marketing and awareness pieces fed by globally aggregated, anonymized telemetry, not by any one organization's own sensors. None of them can show an analyst "what is hitting us right now," because none of them know what is hitting any specific organization.

A third, less technical but equally practical problem concerns communication between SOC staff and the people who decide how to act on what the SOC reports. In under-resourced environments, including much of the open-source, on-premises tooling landscape typical of Azerbaijan's SOC deployments, translating raw event streams into a form non-technical stakeholders can act on frequently does not happen, because analysts do not have the time or dedicated tooling to prepare summaries on top of their primary detection work. A fourth dimension of the problem concerns control over sensitive operational data: once network activity metadata leaves the organization's own infrastructure and is processed by an external vendor, it becomes subject to that vendor's jurisdiction and terms of service, and the resulting dashboards are rarely adjustable to local operational needs.

Taken together, these four problems point toward the same underlying requirement rather than four independent ones. A visualization layer that is driven entirely by an organization's own live alert stream, rather than by an aggregated third-party feed, automatically satisfies the data-sovereignty concern, because nothing needs to leave the network in the first place. A display that is genuinely geographic and animated, rather than tabular, addresses the alert-volume problem directly, since it is precisely the kind of pre-attentive visual encoding the literature associates with faster pattern recognition underload. And a display legible to a non-specialist without additional summarization work closes the communication gap as a side effect of the same design choice, rather than as a separate feature that must be built and maintained on top of it. This convergence is what motivated treating the four problems as a single design target instead of building four separate fixes.

3.2. Problem solution

The system is organized into four tiers: log collection, centralized processing, data enrichment and querying, and visualization delivery. Fig. 1 illustrates the overall architecture of the system.

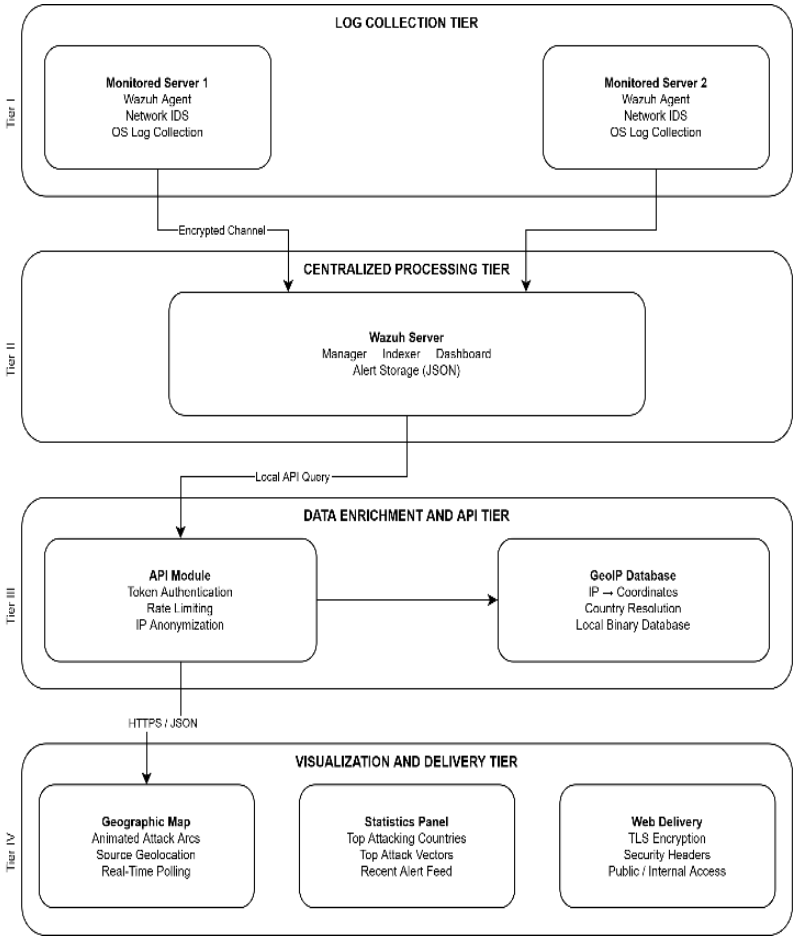

Fig. 1. System Architecture

Log collection. The deployment uses two servers. One handles external web traffic; the other supports backend operations. Both run a Wazuh agent alongside Suricata, an open-source engine that inspects packets against rule sets and raises alerts at the network layer. Agent and Suricata output are merged locally, then forwarded to the central Wazuh Manager.

Centralized processing. The Wazuh Manager receives raw agent data, applies correlation rules, and produces normalized alerts. Wazuh's own architecture separates this into three co-located services: the Manager (rule correlation and alert generation), the Indexer (an OpenSearch-based store exposing an Elasticsearch-compatible API), and the Dashboard (an administrative interface). Alerts are stored as JSON and made queryable through an API endpoint restricted to the local network. Keeping this endpoint unreachable from outside the local network was a deliberate design decision rather than an oversight: the visualization layer described below only ever reads from this endpoint through the intermediate enrichment tier, so the Manager and Indexer themselves are never exposed to the public internet, and a compromise of the public-facing map component cannot by itself grant access to raw alert data or to the underlying detection configuration.

Data enrichment and API. A PHP component sits between alert storage and the visualization front end. For every request it issues three targeted Elasticsearch queries: one for alerts that already carry geolocation metadata, one for Suricata-originated alerts identified by a *data.src_ip* field, and one covering the remainder, tagged instead with *data.srcip*. This three-way split exists because Wazuh and Suricata alerts, although stored in the same index, do not share a single consistent field name for the source address; a naive single-query implementation was tried first and silently dropped a meaningful fraction of Suricata-originated alerts during early testing, which is what motivated separating the queries by source rather than relying on one generic field lookup. Alerts lacking pre-resolved coordinates are enriched on the fly using the MaxMind GeoLite2-City database, a locally stored binary lookup table mapping IP ranges to latitude, longitude, and country. Private, loopback, and link-local ranges are excluded from lookup to avoid placing spurious points on the map. The API

returns a compact JSON payload per alert: coordinates, country, an obfuscated source IP, and the triggering rule's description.

Visualization and delivery. The front end is a single page served over TLS and built with Leaflet.js, an open-source mapping library, using a dark CartoDB basemap chosen for readability in typical SOC lighting conditions. New alerts are drawn as curved, animated arcs from their resolved origin toward a fixed point representing the protected infrastructure. Two ranked panels sit alongside the map, one tracking countries by alert volume, another by triggering rules, and a scrolling feed lists recent events. The last two octets of any source IP shown in this feed are masked, which keeps the interface usable for demonstrations to non-technical audiences without exposing exact addressing. The front end polls the API on a fixed interval rather than holding an open connection, a deliberate simplicity trade-off: a push-based mechanism such as WebSockets would reduce the delay between an alert being stored and it appearing on the map, but polling was chosen for the prototype because it keeps the front end stateless and easy to reason about, at the cost of a small, bounded display latency that has not been an operational problem at current alert volumes. The prototype is deployed at airport.yer.az under an automatically renewed Let's Encrypt certificate and its user interface looks like fig. 2.

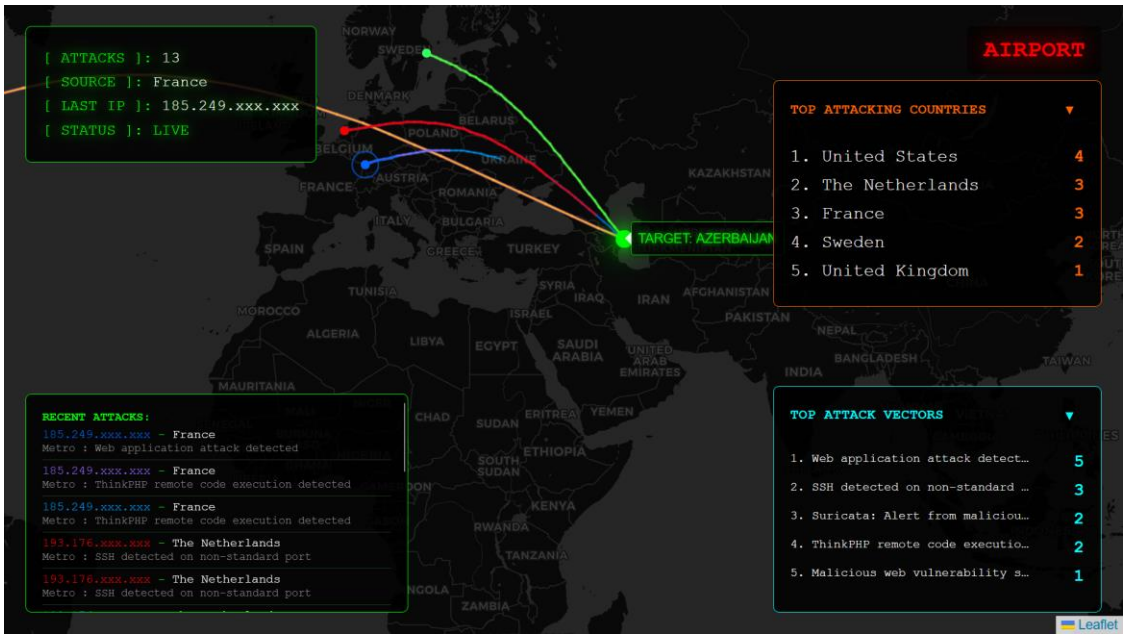

Fig. 2. Attack Map

4. Discussion

To situate the system relative to existing approaches, Table 1 compares it with three commercial threat maps and with Wazuh's own built-in dashboard across five dimensions: whether the data source is organization-specific, whether deployment is on-premises, whether a geographic arc visualization is offered, whether the tool is open source, and whether it depends on any external threat-intelligence feed

Table 1. Comparison of visualization approaches

System	Org-specific data	On-premises	Geo-arc visualization	Open source	No external feed
Kaspersky Cyber Threat Map	No	No	Yes	No	No
Radware Live Threat Map	No	No	Yes	No	No
Check Point ThreatMap	No	No	Yes	No	No
Wazuh Dashboard (built-in)	Yes	Yes	No	Yes	Yes
Proposed system	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**

Commercial maps satisfy the geographic-visualization criterion but fail every other one, since their data is aggregated across many unrelated networks. Wazuh's own dashboard is the mirror image: fully organization-specific and open, but with no geographic arc layer. The proposed system is the only entry satisfying all five criteria simultaneously.

In current operation, the prototype has run continuously against live agent and Suricata output rather than replayed or synthetic data, which is what distinguishes it from most of the visualization research surveyed above. Because the enrichment tier resolves geolocation from locally held databases rather than a cloud lookup service, the pipeline introduces no outbound dependency and continues to function if external connectivity is degraded. Masking the last two IP octets in the public-facing feed has also made it possible to leave the interface running on screens visible to visitors and non-technical staff without exposing exact addressing, which was one of the original communication goals stated above.

The comparison also clarifies what the system deliberately does not attempt to be. It is not a replacement for the correlation and rule-authoring capabilities of the Wazuh Manager, nor for Suricata's packet-level inspection; it is a presentation layer that sits on top of both, in the same spirit as automated visualization pipelines that couple detection and display rather than treating them as separate stages (Alhaidari et al., 2022).

Consequently, the quality of what appears on the map is bounded by the quality of the underlying detections, a deliberate scope decision on the premise that most on-premises SOC deployments already have a working detection stack and are specifically missing the presentation layer described in the problem definition above.

Two practical observations from the deployment period were not anticipated at the design stage and are worth recording separately. The first concerns which part of the interface non-technical viewers actually rely on. During informal demonstrations, visitors and managers tended to glance at the animated map for a general sense of activity level, but formed their actual conclusions, such as "most of this is coming from a particular region" or "one rule is firing unusually often today", from the two ranked panels rather than from the arcs themselves. This suggests that in a dual-audience system of this kind, the summary panels are not a secondary feature attached to the map but carry a comparable share of the communicative load, a distinction the visualization guidelines discussed in Section 2 do not draw out explicitly when they treat "map" and "supporting panel" as a single combined view (Seong et al., 2020).

The second observation concerns cost under concurrency. Because each polling cycle issues three separate Elasticsearch queries per connected client, the API tier's load scales with the number of simultaneously open dashboards as well as with alert volume. This was manageable at the concurrency levels observed during the deployment period described here, but it is the first place contention would be expected to appear if the interface were opened on many more analyst workstations at once, and it is the specific motivation behind the caching and rate-limiting work listed under future directions below. A related limitation is geolocation accuracy itself: the MaxMind GeoLite2-City database used for enrichment is known to misplace address ranges behind carrier-grade NAT, VPN exit nodes, or shared cloud provider blocks. This means that displayed arc should be read as an approximate origin rather than a forensically certified attacker location. This remark applies equally to the commercial threat maps compared in Table 1, none of which publish the accuracy of their own underlying geolocation.

5. Conclusion

This paper described a visualization layer that turns live, on-premises Wazuh and Suricata output into an animated geographic map, built entirely from open-source components and operated inside the organization's own network. Unlike commercial threat maps, it reflects only the organization's own traffic; unlike a standard SIEM dashboard, it renders that traffic geographically and, in a form, readable by non-specialists as well as analysts.

The approach also has limitations. The system currently depends on the completeness and accuracy of the MaxMind GeoLite2-City database, which is known to misplace some IP ranges, particularly those behind carrier-grade NAT, VPN exit nodes, or cloud provider address blocks shared across regions; a geographic arc is therefore a useful approximation of origin, not a certified attribution of an attacker's physical location. The current prototype has been validated through continuous operation against live traffic rather than through a controlled comparison against text-based dashboards, so the cognitive-load and speed-of-comprehension benefits argued for here rest on the

literature cited throughout the paper rather than on a dedicated user study of this specific interface; closing that gap is a priority for future work.

Several extensions are planned. An expanded statistics view would let users break down alert frequency by time period, by rule, and by protocol, and would allow trend comparison across days or weeks rather than only an instantaneous snapshot, which the current top-country and top-rule panels do not support. A drill-down feature would let an analyst open the full raw log behind any point on the map, including source, destination, timestamp, protocol, and payload, without switching tools, closing the gap between the high-level visual and the detailed forensic record an analyst eventually needs. Compatibility with additional detection back-ends, such as other open-source IDS/IPS engines or cloud-native log sources, is planned to broaden applicability beyond the specific stack used in this deployment, alongside continued interface performance work as alert volume and the number of concurrently open dashboards increase. Because the system exposes both a public web page and an internal API, hardening remains a priority: stronger authentication on the API tier, rate limiting to prevent abuse of the enrichment endpoint, and penetration testing before any wider or externally facing deployment. Improving the accuracy of geolocation itself is also a candidate direction, particularly for address ranges known to be poorly served by commercial geolocation databases, such as by combining GeoLite2 with supplementary routing-derived location signals.

Finally, structured empirical evaluation is needed on three fronts: end-to-end latency measurement under realistic and adversarial traffic loads, to confirm that the enrichment and rendering pipeline keeps pace with genuine attack bursts rather than only steady-state traffic; scalability testing at alert volumes substantially higher than those observed during the current deployment period, to establish the operational ceiling of the architecture before it requires horizontal scaling; and structured usability testing with both SOC analysts and non-technical managers, comparing task completion time and situational-awareness accuracy against the text-and-table baseline the system is intended to replace.

Together, these three evaluations would convert the qualitative case made in this paper into quantitative evidence and represent the natural next stage of this research program.

References

Alakbarova, I. (2026). Analysis of the possibilities of applying video analytics in ensuring cyber sovereignty. Problems of Information Society, 17(1), 50–60. https://doi.org/10.25045/jpis.v17.i1.06

Alhaidari, F., Tammas, R., Alghamdi, D., Alrashedi, R., Althani, N., Alsaidan, S., AlFosail, M., Zagrouba, R., & Alattas, H. (2022). A study on automated cyberattacks detection and visualization. In Proceedings of the 14th International Conference on Computational Intelligence and Communication Networks (CICN) (pp. 715–722). IEEE. https://doi.org/10.1109/CICN56167.2022.10008351

Aliyev, R., & Peñalver, L. (2016). Analyzing vulnerability databases. In Proceedings of the 10th IEEE International Conference on Application of Information and Communication Technologies (AICT2016).

Aliyev, R. (2025, December). Establishing a cybersecurity laboratory. In Proceedings of the III International Conference "Sustainable Development Strategy: Global Trends, National Experience and New Goals." Mingachevir State University.

Bhatt, S., Manadhata, P. K., & Zomlot, L. (2014). The operational role of security information and event management systems. IEEE Security & Privacy, 12(5), 35–41. https://doi.org/10.1109/MSP.2014.103

Duo, W., Zhou, M., & Abusorrah, A. (2022). A survey of cyber attacks on cyber physical systems: Recent advances and challenges. IEEE/CAA Journal of Automatica Sinica, 9(5), 784–800. https://doi.org/10.1109/JAS.2022.105548

González-Granadillo, G., González-Zarzosa, S., & Diaz, R. (2021). Security information and event management (SIEM): Analysis, trends, and usage in critical infrastructures. Sensors, 21(14), 4759. https://doi.org/10.3390/s21144759

Jumiaty, & Soewito, B. (2024). SIEM and threat intelligence: Protecting applications with Wazuh and TheHive. International Journal of Advanced Computer Science and Applications, 15(9). https://doi.org/10.14569/IJACSA.2024.0150923

Lallie, H. S., Debattista, K., & Bal, J. (2020). A review of attack graph and attack tree visual syntax in cyber security. Computer Science Review, 35, 100219. https://doi.org/10.1016/j.cosrev.2019.100219

Mahmudova, R. S. (2024). Smart socio-technological infrastructures: Cyber security risks and the human factor. Problems of Information Society, 15(2), 49–60. https://doi.org/10.25045/jpis.v15.i2.06

Seong, Y., Nuamah, J. K., & Yi, S. (2020). Guidelines for cybersecurity visualization design. In Proceedings of the 24th International Database Engineering and Applications Symposium (IDEAS '20) (pp. 1–6). ACM. https://doi.org/10.1145/3410566.3410596

How to cite: Rashad Aliyev, Sarfi Habibova (2026). Real time attack visualization map for security operation centers. Problems of Information Society, 2, 106–112. https://doi.org/10.25045/jpis.v17.i2.12