# Securing Networks with SOAR and SIEM Systems Within Small and Medium Sized Organizations

**Dušan Čatloch, Eva Chovancová, Martin Chovanec, Martin Štancel, Sylvia Maťašová**

Technical University of Kosice, Faculty of Electrical Engineering and Informatics, Dept. of Computers and Informatics, Letná 9, 040 01 Košice, Slovakia, dusan.catloch@tuke.sk, eva.chovancova@tuke.sk, martin.chovanec@tuke.sk, martin.stancel@tuke.sk, sylvia.matasova@tuke.sk

*Abstract: The paper investigates the feasibility of implementing Security Orchestration, Automation, and Response (SOAR) and Security Information and Event Management (SIEM) systems in small to medium-sized organizations. It examines the challenges and benefits of centralized log management, automated security responses, and strategies to balance costs with potential solutions. The research explores different SIEM implementation approaches, comparing on-premises and cloud-based solutions, and emphasizes the importance of automated threat detection through well-defined security alert rules. Additionally, the paper discusses how SOAR systems can enhance incident response capabilities, particularly for organizations with limited or no dedicated security teams. By focusing on these areas, the study highlights the potential for small and medium-sized businesses to adopt robust cybersecurity measures, despite budgetary limitations. It also addresses the integration of these systems into existing IT environments, stressing the importance of customization and scalability to meet the specific needs of smaller enterprises. The paper aims to provide insights into effective methods for protecting diverse and evolving technology infrastructures, ensuring that even resource-constrained organizations can defend against sophisticated cyberthreats efficiently.*

*Keywords: Automation; Orchestration; Security; SIEM; SOAR*

# 1 Introduction

Effective protection of actively used computer networks has always been a challenge. It has been emphasized by the rapid growth and introduction of new technologies and oftentimes also regulatory requirements. For example, the well-known PCI DSS requires tracking of all access to network resources, as well as access to the cardholder data [1]. Not all organizations need to comply with such

requirements, but having access to a well-built centralized logging system is crucial during security incident investigation [2].

The standard solution for solving this problem is the implementation of Security Information and Event Management, also known as SIEM. Software that implements SIEM aims to collect and evaluate logs from various sources, primarily devices of network infrastructure, but also from servers and perhaps sensitive workstations. It allows the employees of security departments to effectively investigate events within a singular console. This introduces time savings [3], but also allows organizations to build rules that can analyze the events and generate alerts automatically [4].

Once the logs are automatically analyzed and alerts are being generated, there is solid base for automated remediation of affected devices, accounts, and other points of contact [5]. Answer to this problem comes with Security Orchestration, Automation, and Response – also known as SOAR. This solution allows automation of commonly needed security-specific actions, such as endpoint isolation, or affected user account lockout [6].

Implementation of these systems, however, tends to be rather costly with licenses reaching tens to hundreds of thousands of Euros annually, which is typically a prohibitive cost for small and medium sized organizations. But there could still be a way to build a well-designed system that can protect organizations of all sizes [7].

The contribution of this paper is in providing a repeatable, resource-conscious deployment methodology for integrating SIEM and SOAR systems in small and medium-sized organizations with restricted budgets and staffing. Existing research predominantly evaluates SIEM or SOAR platforms in isolation or assumes enterprise-grade infrastructure, whereas this paper documents a fully integrated, open-source SIEM–SOAR pipeline based on Elastic Stack, TheHive, and Cortex, including practical deployment architecture, alert-transfer design, and performance behavior under realistic SME constraints. By focusing on cost limitations, architectural scalability, and operational sustainability, the paper fills a gap in current literature, offering actionable guidance for organizations that cannot maintain a full security operations center. The evaluation demonstrates what SMEs can realistically achieve through freely available technologies, highlighting implementation practices that enable meaningful threat detection and automated response without commercial licensing costs.

# 2 Implementation of SIEM Systems

As the introduction suggests, centralized log management has fairly large requirements for data storage [8]. The requirements are specific to the organization within which are deployed and mainly depend on the number of devices logging to

the platform, and data retention requirements. The term centralized logging suggests that all infrastructure devices that have access or process possibly sensitive information should be logging into it.

SIEM is not only about saving data, but also making them accessible, easy to search, and process. For example, Elastic Stack makes use of a custom Elastic Common Schema which describes different data categories and types and sorts them into predictable categories under predictable names. Alongside that a good SIEM system provides visualization utilities that can provide data in human-friendly summaries. As such, they allow to create dashboards for better overview of the protected environment.

## 2.1 Installation Method Considerations

Currently there are two main methods of SIEM installation. There is the traditional method of on-premise installation to servers owned by the organization and the fully managed cloud-based installation. Both of them have their own set of pros and cons which have to be evaluated [9] [10].

The on-premise installation allows the organization to keep full ownership and control over the system. This also means that the organization must have a sufficient hardware infrastructure at its disposal. Given that the organization already has infrastructure and only has to scale it, this may be a viable solution for many [11].

The alternative is a cloud-based fully managed deployment. In this case, the organization only needs to enter a contract with a service provider which will deploy and manage the SIEM. The most tangible downside are the running costs that can grow over time. Unlike the on-premise deployment, costs for just managing the platform will be present, regardless of the system usage.

## 2.2 SIEM Architecture

SIEM systems are designed to be scalable and highly available. This is performed by designing the system as a distributed one. This brings its own set of challenges, but for implementers the biggest one is the requirement of at least three instances of the SIEM. While an administrator can install multiple instances on a single machine, this is inadvisable, as the goal of distributed system is to prevent its failure after loss of a compute node. Because of this reason it is recommended to install SIEM on separate hosts at least, ideally on separate hardware. In the optimal case, there should be at least three hosts on two locations. Offline backup can be achieved by using a NFS storage for data backup, which is different from data replication [12].

The data loses significance over time, and such is also searched less often, which means it can be moved over to the slower, but more cost-efficient data storage.

Within Elastic this is done by using so-called data tiers. There are in total four stages for data within their life cycle: hot, warm, cold, and delete. As described in Fig. 1, incoming data is hot, meaning that they are temporally relevant and are expected to be searched often. After a certain point in time, data is less likely to be searched, and it can be moved to the warm storage. Last stage where the data are still accessible is the cold storage, within which the data can be searched, but the performance will be most likely slow. This tier allows organizations to effectively archive the data for the required time, while still allowing them to look back in the case they need to do so. Once the data are no longer needed, it enters the last stage where it gets closed and deleted off of the cluster.

The implementing engineers must take care to ensure that there are enough replicas with reliable storage and power supply, as well as high-speed interconnections between individual replicas. If there is no other choice, there is nothing preventing them from implementing the entire architecture within one geographical point, of course, but the risk of data loss must be correctly assessed and accepted.

SIEM systems are inherently designed to handle large volumes of data while maintaining their operational integrity. A well-architected SIEM will employ multiple layers of redundancy to ensure that no single point of failure can disrupt operations, a crucial aspect for minimizing downtime. The use of distributed computing allows for higher availability and fault tolerance, ensuring that the system can continue to function even when individual nodes fail. Load balancing mechanisms are typically integrated to distribute data processing tasks across various nodes, preventing bottlenecks and ensuring efficient event ingestion. Backup strategies, including NFS storage, ensure that critical data is safeguarded, while replication across multiple nodes minimizes the risk of data loss. Security remains paramount in the architecture, with robust encryption and access controls governing data flow and storage. Moreover, the ability to integrate machine learning within the architecture enables SIEM systems to detect anomalies and refine alerting rules autonomously, enhancing threat detection accuracy. In small to medium-sized organizations, a well-planned SIEM architecture not only ensures compliance with regulatory standards but also enables efficient resource usage.

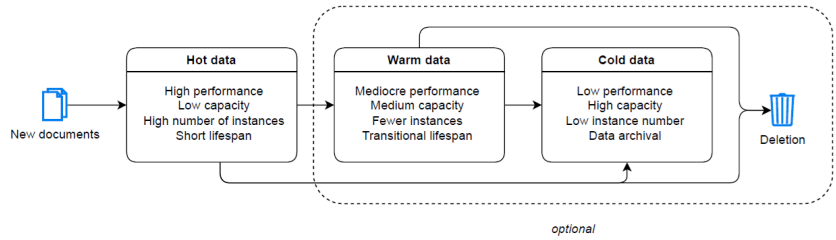


Figure 1
Index lifecycle management progress

## 2.3 Automated Alerting Rules

Data is being ingested to the SIEM all around the clock, but its value can be extracted even when no human operator is at hand. By employing automated detection via security alert rules, administrators can configure the cluster to automatically detect events that have patterns which could be evaluated as malicious [13]. Correctly predicting the specific events and patterns can be complicated, but there are frameworks to assist with the design. Probably the most well-known framework currently is the ATT&CK framework, which describes several stages as well as tactics and techniques that adversaries use during attacks against a protected environment [14] [15].

Not all actions can be detected as single events, so SIEM systems make use of aggregations to identify volumetric attacks or other attacks that can be reliably detected by the count of events as one of the main factors. Additionally, Elastic specifically provides a query language that allows the engineers to describe a series of events that hint on system attack.

Thanks to the open nature of Elastic Stack, it is possible to implement machine learning into the system. While there is ML available within it, it is behind a license barrier, which can be a problem in smaller organizations. There is, however, no license limitation to how much data can be read or processed, excluding the bandwidth and processing power required to move the data. Since the system employs a common scheme for all data processing, it allows for easier development and implementation of Artificial Intelligence models which can be then used to more reliably detect possible security events on the network. They can also be used to sift through the data and detect anomalies, which can be useful to, e.g., detect high-volume data exfiltration, which is known to occur during ransomware attacks [16].

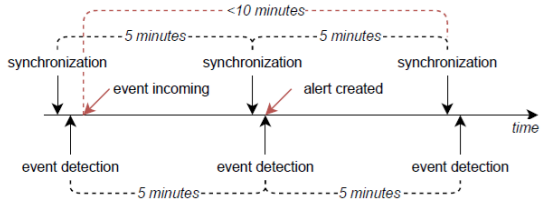


Figure 2
Graphical representation of possible alert delay

The indubitable benefit of automated detection is the possibility of timely detection [17]. All automations can only run so often, so there is an interval during which the events are checked. Suppose that the interval is five minutes and the alerts must be moved to another system afterwards, there is a theoretical longest delay of mere 10 minutes. This interval is illustrated in Fig. 2. In the worst case of about ten minutes an automated action can be triggered, which is arguably faster than a human operator.

# 3 Integration with SOAR Systems

Oftentimes organizations do not have resources available to employ a full scale security operations centre, so they typically assign platform security to the general IT administration teams. Within these teams there are employees with varying levels of experience and knowledge, and sometimes the skills may not be completely up to the task. SOAR aims to at least partially resolve these issues by providing a platform for reponse and remediation [18].

Solutions providing SOAR capabilities are typically also joined with case management system, effectively transforming them into a full-blown Security Incident Response Platforms (SIRP). These platforms can be standalone from SIEM systems, providing a cleaner interface more suited for day to day work. By decoupling they allow source system agnosticism, which means they are able to ingest alerts via a singular interface from many source technologies. One of such systems implementing SIRP functionality is TheHive platform, which will be used as an example within this article [19].

## 3.1 Alert Transfer from SIEM

Because the events are being generated within another system, they must be transferred over to the SIRP [20]. With Elastic Stack this can be done fairly easily by the provided API. TheHive also has an API that receives the events and logs them into the system for further analysis. Connection between these two systems is custom-made for each deployment, and can be done by Python. Usage of this specific programming language is popular amongst system administrators, thanks to its fairly relaxed learning curve, high portability, and easy deployments. You can see the relations in Fig. 3.

*Note: Although the diagram may suggest parallel processing, only the SIEM system performs log ingestion, correlation, and alert generation. The SOAR platform does not analyze raw logs; it consumes exclusively the alerts forwarded by the SIEM. This eliminates analytical redundancy and ensures that detection logic remains centralized within the SIEM.*

During event transfer it is the job of the script and/or its developer to identify which elements are considered observables. These elements are then attached to the alert alongside with the type of the observable. This helps with selection of Analyzers and Responders within the system itself.

To avoid ambiguity regarding the roles of both platforms, it is important to emphasize that the SOAR system receives alerts only from the SIEM and does not independently evaluate incoming logs. SIEM therefore remains the single authoritative source for alert creation, while SOAR provides enrichment, case management, and response orchestration.

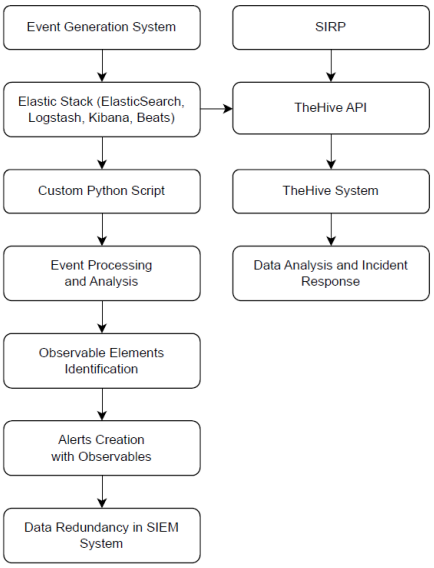


Figure 3
Workflow of Alert Transfer and Incident Response: From Generation to SIEM System Redundancy

Transferring to another system also provides a certain level of redundancy. Should SOAR catastrophically fail, there will still be a copy of data in the SIEM, given that the data was not removed in the meantime or during event transfer. The redundancy applies to alert storage and workflow continuity, but not to duplicate detection logic since only the SIEM evaluates log data for malicious patterns.

## 3.2 System Protection

Many employees, especially at the beginning of their careers, may lack necessary knowledge to protect a network from threats while not causing bigger problems [21].

If the automation contained within SOAR is well implemented, it can prevent from accidental system outage by such internal threats by refusing performing an action that is known to break a system. Such an example may be locking out all access to a specific server [22].

Fig. 4 shows the impact of various factors on cybersecurity in companies. The vertical axis shows the factors that affect cybersecurity [23]. The horizontal axis shows the degree of influence of these factors on a scale of 0 to 10. The first factor is new hires who have a high impact (8.0) on cybersecurity, highlighted in red. This factor is explained as high risk due to lack of experience and lack of knowledge of cybersecurity measures leading to error.

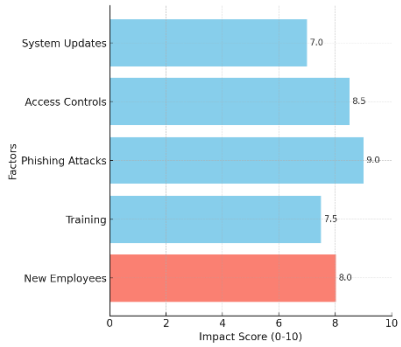


Figure 4

Impact of Various Factors on Cybersecurity in Companies. New Employees have a high impact due to lack of experience and potential errors

The second factor is training with an impact of 7.5, indicating the importance of educating the employees on cybersecurity. The third factor is phishing attacks, which has the highest impact with an impact value of 9.0, highlighting their frequent occurrence and danger to companies. The fourth factor is access controls, which has an impact value of 8.5, showing the criticality of setting access rights correctly. The fifth and final item listed is updating systems, with an impact factor of 7.0, highlighting the importance of regular maintenance and patching of systems.

The protection also extends to automated systems. If automatic response is in place, the system may incorrectly identify benign activity as possibly malicious, and attempt to block it, inadvertently breaking the system. Protection implemented within the automation can also prevent actions in this case since it knows not to perform an action on a protected object.

With implementation of SOAR system the organization can grant very strictly limited permissions to employees that would otherwise be unable to obtain them. The permissions can be granted to the accounts that will be performing the automation while limiting their abilities via the logic of the programs. All usage can be monitored with the built-in audit logging within the application. This way the least privilege principle is retained while allowing security analysts within the organization to perform emergency actions that would otherwise be delegated to other team members.

## 3.3 Additional Data Analysis and Remediative Actions

The design of SOAR systems varies, but TheHive has separated the user platform and the automation platform. This allows for more robust defense-in-depth implementation, as the executive part may be installed on a different system than the user part [24]. The part of the system that performs the actions on behalf of

acting administrators is called Cortex. It is integrated directly with TheHive, which makes it easy to use, even for less skilled employees.

It provides two types of automations: analyzers, and responders. Analyzers take observables and then process them through a specified system, returning the results back to the console for further analysis by the security employees. Based on the collected information the analysts can make better judgements, whether or not to act upon an alert. Responders in turn allow to perform actions on the observables, providing remediation steps necessary to protect the network from further harm [25].

Out of the box there are several connectors provided, but the platform allows implementation of custom programs which can then be registered and used accordingly. There are virtually no limitations, other than the requirement that the scripts must be executable on a Linux-based operating system [26]. Given that the Linux ecosystem is an open one, it can allow for remote execution of programs on different machines and even on different platforms, which makes this whole solution agnostic [27].

## 3.4 Data Privacy

By employing a SIRP solution, an organization has higher level of control over data sharing since they mostly implement the Traffic Light Protocol (TLP) and Permissible Actions Protocol to control the flow of the data within and out of the organization during security event investigation.

The Traffic Light Protocol has been officially published in a guide by NIST [28]. It sets four levels of information security, illustrated by the colors of a traffic light, joined by the white level, later renamed to the clear level. All levels specify with which communities specific information can be revealed and shared.

Permissible Actions Protocol (PAP) has been introduced as a supplement to the TLP. It specifies actions that analysts are able to perform with obtained information. Analogically to the TLP, PAP also employs the colors of a traffic light to signify specified level of protection for specific information [30].

TheHive integrates TLP and PAP deeply into the system, effectively requiring all data to be tagged accordingly. Its safeguards reject action on objects that do not comply with specified configuration, preventing more harm and/or loss of data privacy to the organization or to its customers [29].

# 4 Incident Capture and Processing in SIEM and SOAR Systems

SIEM and SOAR systems play a key role in identifying and processing security incidents in organizations. These systems enable the centralized collection, correlation and analysis of security data from multiple sources, providing a comprehensive, real-time view of the security situation. Once an incident is detected, it is then processed, which includes:

- **Response to the incident**: SOAR systems automate and orchestrate incident response. Response can include isolating the affected system, blocking the malicious IP address, or deploying patches to vulnerable systems.
- **Analysis of the incident**: The security team analyses the incident to understand its origin, scope and impact. This analysis includes a review of the events leading up to the incident and identification of vulnerabilities.
- **Resolving the incident**: After analysis, the incident is dealt with according to established procedures. This may include restoring systems from backups, removing malware and implementing measures to prevent future attacks.
- **Documentation and learning**: Each incident is documented, including details of its detection, analysis and resolution. This information is used to improve future incident responses and to adjust detection policies.

This is shown in the Fig. 5.

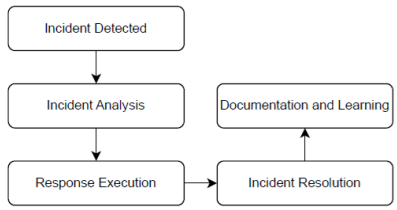


Figure 5
Incident Response

Anomaly detection using TheHive platform focuses on identifying unusual behaviour that may indicate security threats or other problems. When detecting anomalies related to IP addresses that exhibit an excessive number of login attempts, a combination of tools and techniques must be used to effectively monitor and analyze these activities. First, login attempt data is collected from authentication systems such as Active Directory, LDAP, or web servers. These logs contain important information, including IP address, timestamp, type of action (successful or failed login), and user details. Cortex, which is a data analysis and integration tool with TheHive, is used to automatically ingest this data into TheHive. Cortex analyzers allow you to import these logs directly into TheHive.

Rules and alerts are set up in TheHive to detect anomalies. For example, a specific rule could be that if an IP address exceeds a certain number of login attempts in a specified time interval, the system triggers an alert. For example, this rule could be defined such that if an IP address makes more than 10 login attempts in 5 minutes, an alert is automatically triggered.

Cortex is again used to automate the response to detected anomalies, which can perform automated actions such as blocking IP addresses at the firewall, resetting passwords of affected users, or creating an incident for further investigation. Finally, TheHive provides visualization and reporting capabilities to assist in presenting the anomalies found and the management actions taken. With these features, security professionals can better monitor the situation and quickly respond to potential threats [31]. This comprehensive approach allows for effective detection and response to anomalies associated with excessive login attempts from a single IP address, thereby significantly improving infrastructure security. Fig. 6 shows the user interface of the security incident and event management system developed [32].

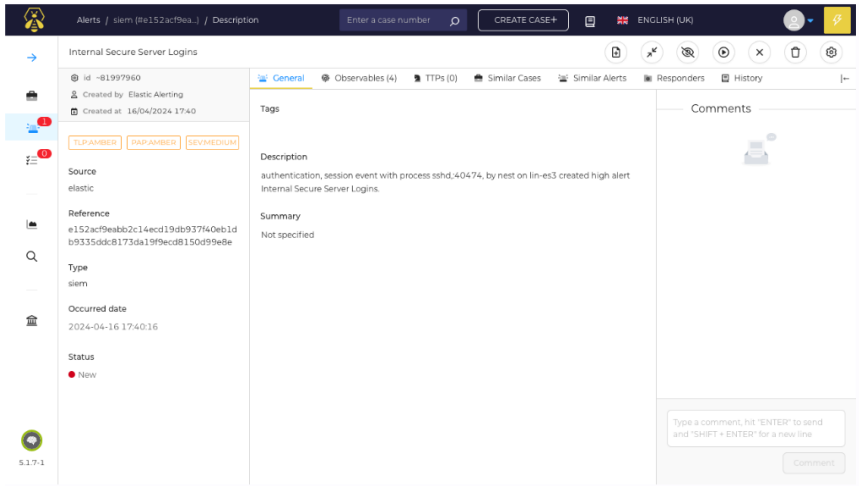


Figure 6
UI

There are several factors that can affect the performance of the SIEM system and cause the system to not process all the events that are sent to it. Some of the most significant factors include hardware resources, misconfiguration, and network issues. Lack of hardware resources is one of the main reasons why such SIEM system may not be able to process all incoming events. If the processor is overloaded, this can lead to delays or the inability of the system to process all events. Similarly, insufficient memory can cause slow processing or complete skipping of some events [33]. Problems can also occur if the system does not have enough disk

space or uses slow disks, which can negatively affect the ability to store and read events efficiently. Incorrect system configuration is another important factor. If a SIEM system is improperly configured, it can lead to problems receiving or processing events. Bugs or errors in the software can cause the system to crash or process events incorrectly, affecting overall system performance. Inefficient or unoptimized code can also cause unnecessary load on the system, resulting in slower event processing. Network issues can significantly affect the performance of a SIEM system. High network latency can cause delays in receiving events, while packet loss can mean that some events never reach the SIEM system [34]. If the network is overloaded, this can lead to delays or loss of events, which directly impacts the system's ability to process all incoming data.

To ensure reproducibility of the performance measurements presented in Sections 4.1 and 4.2, all experiments were executed under controlled conditions. The SIEM platform received log data generated at a constant rate of 1,500 logs per second (LPS) using a synthetic log generator simulating multiple infrastructure components. All tests were performed on a virtual machine equipped with 4 CPUs, 16 GB RAM, and SSD storage, running Elastic Stack version 8.x. Network throughput between log sources and the SIEM was maintained at 1 Gbps, ensuring that no external bottlenecks affected the results. These parameters represent realistic conditions for small and medium-sized organizations.

## 4.1 Overall Processing Success Rate

Measuring the overall success rate of records processed in a SIEM system is key to evaluating the effectiveness of the system. This measurement allows us to see what percentage of the total number of incoming events were successfully processed and how many were not processed for various reasons. With this measurement we can focus and define several metrics that need to be measured:

- Total number of events (Total Events): This number represents the total number of messages that have been sent to the SIEM system. For example, 1000 messages.
- Number of processed events (Processed Events): This number indicates the number of messages that the system has successfully processed. Assume that the system has successfully processed 950 messages.
- Number of failed events (Failed Events): This number represents the number of messages that the system was unable to process. For example, it would be 50 messages (1000 – 950).
- Percentage processing success rate (PPSR): This metric is calculated as the proportion of processed events to the total number of events, expressed as a percentage.

$$PPSR = \left(\frac{Processed\ Events}{Total\ Events}\right) x\ 100 \quad (1)$$

To check the percentage success rate, a measurement was created in which a total of 1,000,000 events were sent to the SIEM system at a stable rate of 1,500 LPS. Of this number, 945,244 events were successfully processed. The remaining events were not processed due to various reasons such as hardware limitations, software bugs or network issues. This means that the system was able to successfully process approximately 94.52% of the total number of events. The remaining 5.48% of events were rejected due to short-term ingestion queue saturation when CPU utilization exceeded approximately 90%. During these peaks, Elastic's internal back-pressure mechanisms temporarily prevented the acceptance of new events to maintain cluster stability. No packet loss occurred at the network layer, and no events were dropped due to disk I/O limitations. While this behavior is expected under resource exhaustion, it indicates that the system reached its processing limits under the tested load. However, organizations with regulatory or audit requirements for complete log retention would need to supplement the architecture with additional compute nodes or an intermediary message-queue buffer to eliminate ingestion loss during peak traffic.

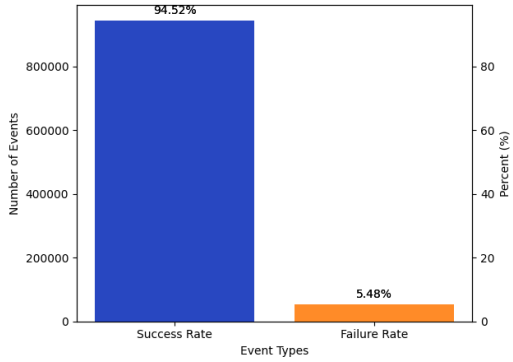


Figure 7
Percentage processing success rate

Fig. 7 shows a graph demonstrating the measured data. The graph shows the number of processed and failed events as well as their percentage. The blue bar represents the number of processed events (945,244) and the gold bar represents the number of failed events (54,756). Above these columns, their percentage values are shown, providing a clear visual overview of the processing efficiency [35]. This metric is crucial for evaluating the performance of the SIEM system because it gives us insight into how efficiently the system is handling incoming data under the specified load conditions (1,500 LPS). A high percentage processing success rate indicates that the system is well configured and powerful enough to handle a large number of events. Conversely, a low percentage success rate may indicate problems that need

to be addressed to improve the efficiency of the system. These issues may include the need to optimize hardware resources, modify the system configuration, or address network issues [36].

## 4.2 Impact of the System Load on Processing Speed

When a system processes a large number of events, its resources are overloaded, which can lead to slower data processing. If the CPU is overloaded, the system may suffer from increased latency, which means that events will be processed more slowly. Similarly, a lack of memory can lead to swapping to disk, which significantly slows down processing. Insufficient disk capacity and speed can also limit the system's ability to store and retrieve data efficiently, which has a direct impact on processing speed.

Fig. 8 shows the relationship between the number of jobs processed per second and the CPU load over time. The incoming event rate during this test was kept constant at 1,500 LPS, ensuring that the throughput decline reflected only the CPU saturation and not changes in input volume. The green line with circular markers shows the number of processed tasks per second, which starts at 140 and gradually decreases. In the interval from 0 to 48 seconds, which corresponds to 0% to 80% of CPU load, the number of processed tasks decreases from 140 to 90. Then, from 48 to 57 seconds, which corresponds to 80% to 95% of CPU load, the number of processed jobs decreases from 90 to 50. Finally, from 57 to 60 seconds, which corresponds to 95% to 100% of CPU load, the number of processed jobs drops from 50 to 30.

This pattern of decreasing number of processed jobs per second depending on CPU load can occur for several reasons. As the CPU load increases, the system has fewer resources available to process additional jobs, resulting in a decrease in the number of jobs that can be processed per second. When the load is high, the system may begin to prioritize some critical tasks over others, which can cause a reduction in the total number of tasks processed. Different processes and tasks may also compete for the same system resources, leading to increased latency and lower processing efficiency.

When the system is fully loaded, i.e., 100% CPU load, several situations can occur. The system no longer has free resources to process new jobs, leading to a dramatic drop in performance. The time required to process jobs increases significantly, which can affect the time sensitivity of applications and services. If the system continues to accept jobs beyond its capacity, this can lead to congestion and possible outages where some jobs will not be processed at all.

Together, these factors can significantly affect the performance of a SIEM system. To minimize these issues, it is important to ensure sufficient hardware resources, proper system configuration, and a robust network infrastructure [37]. Regular maintenance and software optimization are also key to ensuring that the system will

be able to efficiently process all incoming events and provide accurate and timely information for network security.

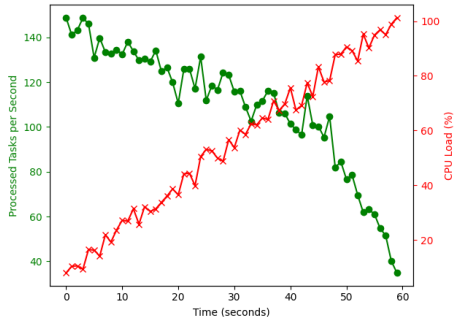


Figure 8
Impact of the system loadin on a processing speed

# 5 Financial Considerations

A small to mid-sized organization typically operates with strict financial constraints, making the adoption of advanced security platforms challenging. A detailed cost comparison highlights differences between on-premises and cloud-based SIEM and SOAR deployments. An on-premises Elastic Stack deployment similar to the environment used in this study requires an initial hardware investment of approximately 2,000–3,000 EUR for server-grade equipment, with yearly operating costs of 200–400 EUR for electricity and maintenance. In contrast, commercial cloud-based SIEM platforms commonly charge between 0.20 and 0.35 EUR per ingested gigabyte, resulting in annual expenses exceeding 6,000 EUR for organizations generating only 50–70 GB of logs per day.

SOAR systems show comparable disparities. TheHive and Cortex can be deployed without licensing fees, while commercial SOAR platforms typically start at 20,000 EUR per year, placing them outside the reach of many SMEs. For this reason, open-source tools represent a financially viable option that enables organizations to adopt automated response capabilities that would otherwise be unaffordable. Using free SIEM/SOAR platforms also allows SMEs to allocate budget toward hardware scaling rather than software licensing, improving processing reliability under peak loads.

Early integration of SIEM and SOAR provides long-term financial advantages. As organizations grow, regulatory requirements and audit obligations become more stringent. Having a scalable, open-source security stack in place reduces the cost of achieving compliance, eliminating the need for disruptive and expensive late-stage implementations.

## Conclusion

The successful implementation of SIEM and SOAR systems in small and medium-sized organizations is not only feasible but essential for establishing a comprehensive security framework in today's evolving threat landscape. Although the upfront and operational costs, particularly for infrastructure and hardware, might be daunting, the long-term benefits far outweigh these challenges. By centralizing log management, organizations can efficiently monitor and correlate security events, gaining invaluable insights into potential vulnerabilities and attack vectors. Moreover, SIEM systems, when integrated with automated alerting mechanisms, enable faster and more accurate threat detection, reducing the time it takes to respond to incidents.

For small and medium-sized enterprises with limited IT resources, coupling SIEM with SOAR provides an even greater advantage. SOAR systems automate critical security responses – such as isolating compromised systems or locking user accounts – without requiring immediate human intervention. This automation is particularly vital for organizations that cannot afford dedicated security teams, as it mitigates the impact of cyberattacks while simultaneously relieving the workload on general IT staff. Moreover, SOAR's incident response capabilities allow organizations to investigate and remediate issues in a structured manner, ensuring that no threat is overlooked, and that post-incident analysis can inform better defenses in the future.

Furthermore, small and medium-sized enterprises stand to benefit from the scalability of SIEM and SOAR solutions. As businesses grow, their security needs will evolve, and these systems are designed to scale with them. Cloud-based solutions offer flexibility and cost-efficiency, while on-premises systems provide greater control and customization. In either case, having a solid SIEM architecture in place early on allows businesses to prepare for audits and regulatory requirements, ensuring that compliance is achieved without the need for disruptive overhauls later.

Another key takeaway is the value of integrating machine learning and artificial intelligence within SIEM systems. This capability can significantly enhance threat detection by identifying patterns and anomalies that would be missed through manual processes alone. For example, machine learning models can be trained to spot unusual network behavior, flagging potential threats like data exfiltration or distributed denial-of-service attacks before they escalate. By leveraging these technologies, small and medium-sized enterprises can protect themselves against sophisticated cyberthreats that traditionally only larger organizations could fend off.

Ultimately, the combination of SIEM and SOAR systems offers a robust, scalable, and efficient solution for improving network security. Even organizations constrained by budget or personnel limitations can adopt these technologies to not only defend against increasingly sophisticated cyberthreats but also to ensure long-term resilience and growth. Investing in these solutions early lays the groundwork

for a sustainable security posture, ensuring that small and medium-sized enterprises remain competitive, compliant, and secure in an increasingly digital world.

**References**

[1] Ataya G.: Pci dss audit and compliance, Information security technical report, Vol. 15, No. 4, pp. 138-144, 2010

[2] Lee J., Kim J., Kim I. and Han K.: Cyber Threat Detection Based on Artificial Neural Networks Using Event Profiles, IEEE Access, Vol. 7, pp. 165 607-165 626, 2019

[3] Rinnan R.: Benefits of Centralized Log File Correlation. Master's Thesis, Norwegian University of Science and Technology (NTNU), 2005

[4] Podzins O. and Romanovs A.: Why SIEM is Irreplaceable in a Secure IT Environment? in 2019 Open Conference of Electrical, Electronic and Information Sciences (eStream) IEEE, pp. 1-5, Apr. 2019

[5] Kovacevic B.: Security Orchestration, Automation, and Response for Security Analysts. pp. ISBN: 978-1803242910, 2023

[6] Wahab O., Mourad A., Otrok H. and Taleb T.: Federated machine learning: Survey, multi-level classification, desirable criteria and future directions in communication and networking systems. IEEE Commun. Surveys Tuts., Vol. 23, No. 2, pp. 1342-1397, 2nd Quart. 2021

[7] Fernandes D. A., Soares L. F., Gomes J. V., Freire M. M. and Inacio P. R.: A Quick Perspective on the Current State in Cybersecurity. Emerging Trends in ICT Security, pp. 423-442, 2014

[8] Lopez J. M.: Systematic review of SIEM technology: SIEMSC birth. International Journal of Information Security, No. 22, p. 691-711, Jan. 2023

[9] Vikas S., Gurudatt K., Vishnu M. and Prashant K.: Private vs public cloud. International Journal of Computer Science and Communication Networks, Vol. 3, No. 2, p. 79, 2013

[10] Hussein M. K., Bin Zainal N. and Jaber A. N.: Data security analysis for DDoS defense of cloud based networks, in Proc. IEEE Student Conf. Res. Develop. (SCOReD), Kuala Lumpur, Malaysia, pp. 305-310, Dec. 2015

[11] Hooper E.: Intelligent Techniques for Network Sensor Information Processing in Large-Scale Network Infrastructures, International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP 08), IEEE CS, pp. 593-598, 2008

[12] Qian C., Zhao X. Yang G. and Nakagawa T.: Optimal Backup Policies for A Database System with Full and Dierential Backup, In S.Sheu, T.Do(ED), Advanced Reliability Modeling III, McGraw Hill, pp. 663-669, 2010

[13] Wu K., Chen Z. and Li W.: A novel intrusion detection model for a massive network using convolutional neural networks, IEEE Access, Vol. 6, pp. 50850-50859, 2018

[14] Strom B. E., Applebaum A., Miller D. P., Nickels K. C., Pennington A. G. and Thomas C. B.: Mitre attack: Design and philosophy, in Technical report. The MITRE Corporation, 2018

[15] Gamer T.: Anomaly-Based Identification of Large-Scale Attacks, in Global Telecommunications Conference. GLOBECOM 2009. IEEE, pp. 1-6, 2009

[16] Ullah F., Edwards M., Ramdhany R., Chitchyan R., Babar A. and Rashid A.: Data exfiltration: A review of external attack vectors and countermeasures, Journal of Network and Computer Applications, Vol. 101, pp. 18-54, 2018

[17] Vinayakumar R., Alazab M., Soman K., Poornachandran P., Al-Nemrat A. and Venkatraman S.: Deep learning approach for intelligent intrusion detection system, IEEE Access, Vol. 7, pp. 41525-41550, 2019

[18] Brewer R.: Could soar save skills-short socs? Computer Fraud and Security, Vol. 2019, No. 10, pp. 8-11, 2019

[19] OpenSOC: Open-Source SIEM and SOAR Architecture for Real-Time Cybersecurity Operations. GitHub, 2021

[20] Schlette, D., Caselli, M. and Pernul, G. (2021) A comparative study on cyber threat intelligence: The security incident response perspective. IEEE Communications Surveys & Tutorials, 23(4), pp. 2525-2556

[21] Dionísio N., Alves F., Ferreira P. M. and Bessani A.: Cyberthreat detection from Twitter using deep neural networks, in Proc. IEEE Int. Joint Conf. Neural Netw. (IJCNN), pp. 1-8, 2019

[22] Zhu Z. and Dumitras T.: Chainsmith: Automatically learning the semantics of malicious campaigns by mining threat intelligence reports, in Proc. IEEE Eur. Symp. Security Privacy (Euro S&P), pp. 458-472, 2018

[23] Zhong C., Yen J., Liu P. and Erbacher R. F.: Automate Cybersecurity Data Triage by Leveraging Human Analysts' Cognitive Process, in 2016 IEEE 2nd International Conference on Big Data Security on Cloud (BigDataSecurity), IEEE International Conference on High Performance and Smart Computing (HPSC), and IEEE International Conference on Intelligent Data and Security (IDS). IEEE, pp. 357-363, Apr. 2016

[24] Kuipers D. and Fabro M.: Control systems cyber security: Defense in depth strategies, Idaho National Lab. (INL), Idaho Falls, ID (United States), Tech. Rep., 2006

[25] Hochreiter S., Bengio Y., Frasconi P., Schmidhuber J., Kolen J. F. and Kremer S.: Gradient flow in recurrent nets: The difficulty of learning long-term dependencies, in A Field Guide to Dynamical Recurrent Networks. Piscataway, NJ, USA: IEEE Press, pp. 237-243, 2001

[26] Roig J. S. P., Gutierrez-Estevez D. M. and Gündüz D.: Management and orchestration of virtual network functions via deep reinforcement learning, IEEE J. Sel. Areas Commun., Vol. 38, No. 2, pp. 304-317, Feb. 2020

[27] Susukailo V. and Lakh Y.: Access control system based on encryption in QR-Code technology, IEEE 4th International symposium on wireless systems within the international conferences on intelligent data acquisition and advanced computing systems (IDAACS-SWS 2018): Proceedings, Lviv, pp. 158-161, Sep. 2018, DOI: https://doi.org/10.1109/IDAACS-SWS.2018.8525779

[28] Johnson C., Badger L., Waltermine D., Snyder J., Skorupka C. et al: Guide to cyber threat information sharing, NIST special publication, Vol. 800, No. 150, p. 35, 2016

[29] Behzadan V., Aguirre C., Bose A. and Hsu W.: Corpus and deep learning classifier for collection of cyber threat indicators in Twitter stream, in Proc. IEEE Int. Conf. Big Data (Big Data), pp. 5002-5007, 2018

[30] CERT-FR: Anssi policy for sharing and handling its operational information [Online, 2024] Available: cert.ssi.gouv.fr/csirt/sharing-policy/

[31] Benjamin V., Li W., Holt T. and Chen H.: Exploring threats and vulnerabilities in hacker Web: Forums, IRC and carding shops, in Proc. IEEE Int. Conf. Intell. Security Inf. (ISI), pp. 85-90, 2015

[32] Naseer .S, Saleem Y., Khalid S., Bashir M. K., Han J., Iqbal M. M. and Han K.: Enhanced network anomaly detection based on deep neural networks, IEEE Access, Vol. 6, pp. 48231-48246, 2018

[33] Horzyk A., Starzyk J. A. and Graham J.: Integration of Semantic and Episodic Memories, IEEE Transactions on Neural Networks and Learning Systems, Volume 28, Issue: 12, 2017

[34] Musarurwa S., Gamundani A. M. and Shava F. B.: A review of security challenges for control of access to Wi-Fi networks in tertiary institutions, In 2017 IST-Africa Week Conference (IST-Africa) IEEE, pp. 1-8, 2017

[35] Sun N., Zhang J., Gao S. Zhang L. Y., Camtepe S. and Xiang Y.: Cyber information retrieval through pragmatics understanding and visualization, in IEEE Transactions on Dependable and Secure Computing, Vol. 20, No. 2, pp. 1186-1199, Mar. 2023

[36] Hu W., Hu W. and Maybank S.: AdaBoost-based algorithm for network intrusion detection, IEEE Trans. Syst., Man, Cybern. B, Cybern., Vol. 38, No. 2, pp. 577-583, Apr. 2008

[37] Kotenko I., Polubelova O. and Saenko I.: The Ontological Approach for SIEM Data Repository Implementation, Green Computing and Communications (GreenCom), 2012 IEEE International Conference, pp. 761-766, Nov. 2012