**Computer Fraud and Security**
**ISSN (online): 1873-7056**

# Insider Threats and Information Security Risks in Modern Enterprise Network Environments

**Partha Chakraborty**

Department of Computer Science, Justice Abu Zafar Siddiquie Science and Technology University,Kumurtuli,Bangladesh

Email: parthachk64@gmail.com

https://orcid.org/0009-0006-3203-8902

**Abstract**

Modern enterprise networks have changed significantly as organizations have adopted distributed workforces, hybrid cloud environments, and more flexible access models. These changes have weakened the effectiveness of traditional perimeter-based security, where organizations largely relied on keeping threats outside the network. Today, identity and access have become much more important, and insiders—individuals who already have legitimate credentials and authorized access—represent a major information security risk. This paper examines insider threats and information security risks in modern enterprise environments through a structured review of peer-reviewed literature. It first establishes the basic concepts of insider threats by distinguishing between malicious, negligent, and accidental insiders and reviewing established classification models. It then examines common insider threat types and motivations, including data thieves, saboteurs, spies, and other malicious insider profiles. The paper also discusses information security risks associated with modern enterprise networks and evaluates different detection approaches, including rule-based and behavior-based methods and machine learning models such as autoencoders, Isolation Forest, and XGBoost. Preventive approaches, including Zero Trust, privileged access management, just-in-time access, and data loss prevention, are also examined. Based on these findings, the paper proposes an integrated defense-in-depth framework that combines technical detection, access governance, and human oversight. Finally, it discusses ongoing challenges such as class imbalance, limited ground truth, behavioral changes, evasion techniques, and employee privacy, along with potential directions for future research.

**Keywords:** insider threat; information security; enterprise networks; user behavior analytics; machine learning; Zero Trust; privileged access management; data loss prevention

## 1. Introduction

Modern enterprise network environments look very different from the smaller, more contained networks that many traditional security architectures were originally designed to protect. Distributed workforces, hybrid cloud infrastructure, software-as-a-service platforms, and mobile devices have gradually broken down the traditional network perimeter. As a result, identity has become a much more important security control than physical network location. In this environment, insiders can cause significant damage because they already have legitimate access to organizational systems and information and are often trusted by the organization (Renaud et al.,). Unlike an outside attacker, an insider does not necessarily need to break through a security boundary. Their credentials, privileges, and knowledge of the organization may already give them everything they need to cause harm.

This makes insider threats particularly difficult and potentially damaging. Malicious insider attacks are considered one of the most serious threats to the networked systems of companies and government agencies (Le et al., 2020). Insider threats can result in major financial, operational, and reputational consequences, while traditional security tools such as firewalls and antivirus software may not be sophisticated enough to identify and stop these activities in real time (Qawasmeh & AlQahtani,). More broadly, insider threats and data breaches continue to create serious challenges for modern organizations (Thapaliya & Gurung,).

The limitations of traditional security controls are not simply technical oversights. They are built into the way these systems operate. Traditional data loss prevention systems often depend on static rules and keyword-based detection, which makes them less effective against changing insider behavior. These systems may miss subtle behavioral changes or attempts to hide data exfiltration, while also generating large numbers of false positives (Thapaliya & Gurung,). Since insiders already have legitimate access, their suspicious behavior may develop gradually rather than appearing as an obvious security violation. This makes continuous monitoring and long-term behavioral analysis increasingly important (Ayobami et al.,).

This paper provides a structured review of insider threats and information security risks in modern enterprise network environments. It has four main objectives: (i) to establish the conceptual foundations and classifications of insider threats; (ii) to examine common insider threat types and motivations; (iii) to analyze the information security risks associated with

modern enterprise networks; and (iv) to evaluate detection and prevention strategies and bring them together into an integrated framework.

**2. Conceptual Foundations: Defining the Insider Threat**

Insider threats involve harmful or malicious actions carried out by individuals who have authorized access within an organization. They represent a major cybersecurity concern because an insider can potentially cause more damage than an external attacker due to the access and knowledge they already possess (Al-Shehari & Alsowail, 2021). The important point is that an insider threat is defined less by the technical method used and more by the legitimacy of the access involved. An insider may not need to bypass authentication, break network segmentation, or exploit a software vulnerability because the necessary access has already been provided.

The literature generally divides insider activity into three broad categories. Based on Sinclair and Smith's classification, one widely used framework distinguishes between individuals who intentionally carry out malicious actions, individuals who act for personal financial gain, and individuals who unintentionally cause harm through accidents or mistakes (Homoliak et al., 2019). This classification is useful because it separates intent from impact. An employee does not necessarily have to be malicious to cause significant damage. Negligence, mistakes, or cooperation with an outside attacker can also result in serious security incidents.

This distinction has important implications for mitigation. Organizations need to understand the different types of insider threats before they can determine the appropriate response to each one (Renaud et al.,). This idea contributed to the development of VISTA (inclusive Insider Threat tAxonomy), which was developed through an extensive literature review and a survey of C-suite executives. The goal was to create a taxonomy that was both academically grounded and useful to organizational leaders (Renaud et al.,).

Another important part of the insider threat problem is the human element. Employees who intentionally want to harm their organizations present a serious security risk because they already have access to sensitive information and are generally trusted by the organization (Weber et al., 2020). Trust can therefore become a vulnerability when it is not supported by continuous verification. This helps explain the movement away from static, trust-based access toward risk-based access that is continuously evaluated.

**3. Typologies and Motivations of Insider Threats**

Understanding why insiders act is important because motivation can influence both behavior and the appropriate response. Weber et al. identify six types of malicious insiders based on real-world cases: Whistleblowers, Spies, Scammers, Saboteurs, Malicious Enablers, and Data Thieves. Each type can arise from different personal motivations or external circumstances (Weber et al., 2020). Their research also emphasizes the value of identifying warning signs early because intervention may prevent an employee from becoming malicious in the first place (Weber et al., 2020).

These warning signs can be both psychological and situational. For example, personal characteristics such as greed, introversion, or difficulty accepting criticism may combine with stressors such as frustration or dissatisfaction. Changes in behavior, including unusual working hours or unexpected travel, may also provide additional warning signs (Weber et al., 2020). The important point is that one indicator by itself may not mean much. It is the combination of multiple signals that can provide a stronger indication of potential risk. This is one reason insider threat programs need to correlate information from multiple sources rather than rely on individual alerts.

Homoliak et al. provide another detailed classification of insider motivations. Their work identifies several archetypes. Machiavellians may engage in sabotage, industrial espionage, intellectual property theft, or other actions that could help advance their careers. Proprietors may develop a strong sense of ownership over the systems they manage and may attempt to destroy systems rather than surrender control. Avengers are disgruntled individuals who may act impulsively because they believe they have been treated unfairly. Career thieves may enter an organization with the intention of committing fraud or stealing intellectual property. Moles may seek to steal intellectual property for a foreign government, organization, or competing company (Homoliak et al., 2019). The same research also notes that insiders may install logic bombs that can act as a form of employment insurance if their activities are discovered (Homoliak et al., 2019).

Taken together, these classifications show that insider threats are not one single type of problem. Different insiders have different motivations and therefore require different responses. Career thieves and moles who begin employment with malicious intentions may not be strongly influenced by security awareness training. Disgruntled employees may be better addressed through grievance processes and management intervention, while negligent employees may benefit more from technical safeguards and better system design. Renaud et al.'s work is particularly useful because it connects insider threat categories with specific mitigation strategies (Renaud et al.,).

## 4. Information Security Risks in Modern Enterprise Network Environments

### 4.1 Data Exfiltration and the Departing Insider

One of the most concerning insider scenarios involves an employee stealing or leaking information shortly before leaving an organization (Al-Shehari & Alsowail, 2021). This period can be especially risky because the employee may still have full access to systems while already having decided to leave. Data can also be removed gradually over time, making the activity appear normal at first. Because of this, researchers have paid particular attention to identifying suspicious behavior during the period leading up to an employee's departure (Al-Shehari & Alsowail, 2021).

### 4.2 Evasion of Content-Based Controls

Content-based security controls can be especially vulnerable when insiders understand how those systems work. Employees with access to confidential information can modify the wording of documents by using synonyms, potentially bypassing DLP systems that depend heavily on keyword matching (Montano et al.,). DLP systems typically monitor information being transferred or exfiltrated by comparing the sensitivity of newly generated information with known sensitive information (Montano et al.,).

The weakness of keyword-based systems becomes clear when considering how easily wording can be changed. An enterprise anti-leakage prototype, for example, uses keyword-based content monitoring to identify confidential information in documents before they are sent. When sensitive keywords are detected, both the user and network administrator can be alerted and the event can be recorded in a log (Montano et al.,). Other approaches have explored different forms of behavioral information, including eye-tracking technology, which can capture information such as gaze location, gaze patterns, and points of interest (Montano et al.,).

### 4.3 Encryption, Telemetry Gaps, and the Limits of Content Inspection

Another challenge is that content inspection becomes less useful when data is encrypted. In one documented insider case, an employee encrypted outbound data packets, making traditional DLP tools ineffective. Instead, the detection system relied on endpoint telemetry and metadata, including filename entropy and file movement frequency, to identify signs of possible data exfiltration (Ayobami et al.,).

This example reflects a broader change in cybersecurity. Encryption is now a normal part of legitimate enterprise communication, which means security systems cannot always rely on examining the actual contents of network traffic. As a result, behavioral information and metadata are becoming increasingly important for detecting suspicious activity.

### 4.4 Identity, Privilege, and the Dissolution of the Perimeter

The combination of Zero Trust and Privileged Access Management represents an important shift away from traditional perimeter-based security toward identity-centered security. These approaches continuously verify users regardless of their location or privilege level (Syed, ). Micro-segmentation helps contain potential breaches, while least-privilege principles reduce the amount of access available to each user (Syed,).

Remote work and hybrid cloud adoption have made the weaknesses of traditional perimeter security even more obvious. Privileged accounts are especially important because they can provide broad access across an enterprise environment (Subramanyam,).

Privilege accumulation can also create significant risks. Traditional PAM approaches may leave organizations vulnerable to "privilege creep" and unmonitored "orphaned accounts" (Subramanyam,). Orphaned accounts are particularly concerning because credentials can remain active after an employee changes roles or leaves the organization. This creates a situation where legitimate access continues to exist even though the original justification for that access may no longer be present.

## 5. Detection: From Rules to Behavioral Analytics

### 5.1 The Limits of Rule-Based and Signature-Based Detection

Rule-based and signature-based systems assume that malicious activity can be identified in advance and matched against known patterns. This approach is difficult to apply to insider threats for several reasons. First, insiders often use legitimate tools and permissions, meaning there may be no obvious malicious signature. Second, insiders can adapt their behavior to avoid detection. For example, an employee with access to confidential information can deliberately change the wording of files to avoid keyword-based detection (Montano et al.,).

Static rules therefore struggle to detect subtle behavioral changes and obfuscated data exfiltration. They can also generate high numbers of false positives while still missing genuine malicious activity (Thapaliya & Gurung, ).

### 5.2 Behavioral Baselining and User Behavior Analytics

Behavioral analytics takes a different approach. Instead of asking whether a specific action is always prohibited, it establishes what normal behavior looks like for a particular user and then looks for meaningful deviations.

One example involves an employee who had maintained a relatively consistent access pattern during eight years with a company. During the employee's final week, however, they began accessing sensitive research and development folders between 2 and 4 AM rather than during their normal 9–5 schedule. The system identified the unusual access times through a user behavior analytics engine. The late-night access became more concerning when it was combined with other changes: the employee began compressing multiple high-value files and transferring them to a personal cloud storage account, something they had never done before. Machine learning models compared this new behavior with the employee's historical activity as well as the behavior of peer researchers who followed appropriate data-handling practices (Ayobami et al.,).

What made the situation suspicious was not one prohibited action. Instead, it was the combination of several activities that represented a clear departure from the employee's normal behavior.

Behavioral analytics has also been expanded through natural language processing. User and Entity Behavior Analytics can establish behavioral baselines and identify unusual activity, while NLP can analyze unstructured information such as emails, chat messages, and documents using semantic analysis, sentiment detection, and entity recognition (Thapaliya & Gurung,). Combining these approaches can help organizations identify both who is behaving unusually and what information may be at risk.

### 5.3 Machine Learning Models and Reported Performance

Research on machine learning-based insider threat detection has produced promising results across several different model types. Le et al. developed a user-centered detection system that analyzes behavior at multiple levels. Their work recognizes several major challenges, including highly imbalanced data, limited ground truth, and changes in user behavior over time (Le et al., 2020). Their evaluation found that the system could detect up to 85% of malicious insiders while maintaining a false positive rate of only 0.78%. The system could also detect suspicious behavior as little as 14 minutes after the first malicious action (Le et al., 2020).

Other studies have compared different machine learning algorithms. Ibraheem et al. evaluated Isolation Forest, One-Class SVM, and Autoencoder models using a comprehensive user activity dataset. Their results showed that the Autoencoder performed better than the other models based on precision, recall, F1-score, and ROC-AUC measures (Ibraheem et al.,). These findings suggest that combining anomaly-based user behavior analytics with machine learning can improve the ability to identify insider threats proactively.

Research has also focused specifically on detecting data leakage before employees leave an organization. Al-Shehari and Alsowail addressed potential bias related to data encoding by using feature scaling and one-hot encoding. They also addressed class imbalance using the synthetic minority oversampling technique (Al-Shehari & Alsowail, 2021). When applied to the CMU-CERT Insider Threat Dataset and compared with ground truth, their model achieved an AUC-ROC value of 0.99, outperforming other approaches evaluated on the same dataset (Al-Shehari & Alsowail, 2021).

More recent research has focused on real-time classification and risk levels. Qawasmeh and AlQahtani developed a machine learning system that combines real-time anomaly detection with dynamic user profiling. The system categorizes employees into low-, medium-, and high-risk groups. Using a synthetic dataset, the system achieved very high accuracy across several machine learning models, with XGBoost performing the best (Qawasmeh & AlQahtani,). Risk classification can be particularly useful in practice because security teams can focus their limited investigative resources on the highest-risk situations instead of treating every alert equally.

## 6. Preventive Architecture: Zero Trust, Privileged Access Management, and Data Loss Prevention

Detecting insider threats is important, but detection alone is not enough. Organizations also need to limit the amount of damage an insider can cause in the first place. Three major approaches are particularly important: Zero Trust, privileged access management, and contextual data loss prevention.

**Zero Trust and identity-centered access.** Zero Trust combined with PAM requires continuous verification of users regardless of where they are located or what level of privilege they have (Syed,). Micro-segmentation can help contain breaches, while least-privilege principles reduce unnecessary access. Technologies such as passwordless authentication, behavioral analytics, and cloud-native security solutions have made Zero Trust architectures increasingly practical for privileged access. Organizations using these approaches have demonstrated stronger security postures, reduced breach impacts, faster threat detection, and improved operational efficiency (Syed,).

**Just-in-time privileged access.** One of the biggest weaknesses in traditional privileged access management is standing privilege. An always-active administrative account is essentially an always-available target. Just-in-time access addresses this problem by providing temporary, time-limited privileges to human and nonhuman identities only when they are

actually needed. This reduces the exposure created by standing privileges and helps address privilege creep and orphaned accounts (Subramanyam,). According to Subramanyam, JIT access can reduce the threat window associated with privilege attacks by more than 90%. Risk-based authentication and self-service workflows can also help organizations maintain strong security without creating unnecessary barriers for employees. The broader goal is to move from "trust-by-default" toward "trust-by-exception" (Subramanyam,).

**Data loss prevention with contextual intelligence.** Keyword-based DLP can be bypassed through simple changes in wording, such as replacing sensitive terms with synonyms (Montano et al.,). Modern DLP therefore needs to consider behavioral and semantic context in addition to document content. Combining behavioral analytics with NLP can improve the ability of DLP systems to identify both suspicious users and sensitive information in real time (Thapaliya & Gurung,).

**7. An Integrated Framework for Insider Risk Management**

The research reviewed in this paper suggests that no single security control is sufficient to manage insider threats effectively. Instead, organizations need multiple layers that complement one another. An integrated framework for modern enterprise environments can be organized into four layers.

**Layer 1 — Identity and access governance.** Organizations should enforce least privilege and continuous verification, use just-in-time access to reduce standing privileges, and regularly address orphaned accounts and privilege creep (Syed,)(Subramanyam,).

**Layer 2 — Continuous behavioral detection.** Organizations should establish individualized and peer-based behavioral baselines using UEBA and apply machine learning models such as Autoencoders, Isolation Forest, One-Class SVM, and gradient-boosted classifiers to identify unusual behavior in real time (Ayobami et al.,)(Ibraheem et al.,)(Qawasmeh & AlQahtani,). Class imbalance and encoding bias also need to be addressed because they can significantly affect detection accuracy (Al-Shehari & Alsowail, 2021).

**Layer 3 — Contextual data protection.** Content inspection should be combined with semantic NLP analysis and metadata-based telemetry, including information such as filename entropy and file movement frequency. This provides additional visibility when outbound data is encrypted and cannot be inspected directly (Thapaliya & Gurung,)(Ayobami et al.,).

**Layer 4 — Human review, response, and governance.** Machine-generated risk scores should support, rather than completely replace, human decision-making. Low-, medium-, and high-risk classifications can help security teams prioritize investigations (Qawasmeh & AlQahtani,). At the same time, behavioral warning signs such as unusual working hours, dissatisfaction, or sudden changes in file-handling behavior can provide opportunities for early intervention before an employee becomes actively malicious (Weber et al., 2020)(Ayobami et al.,).

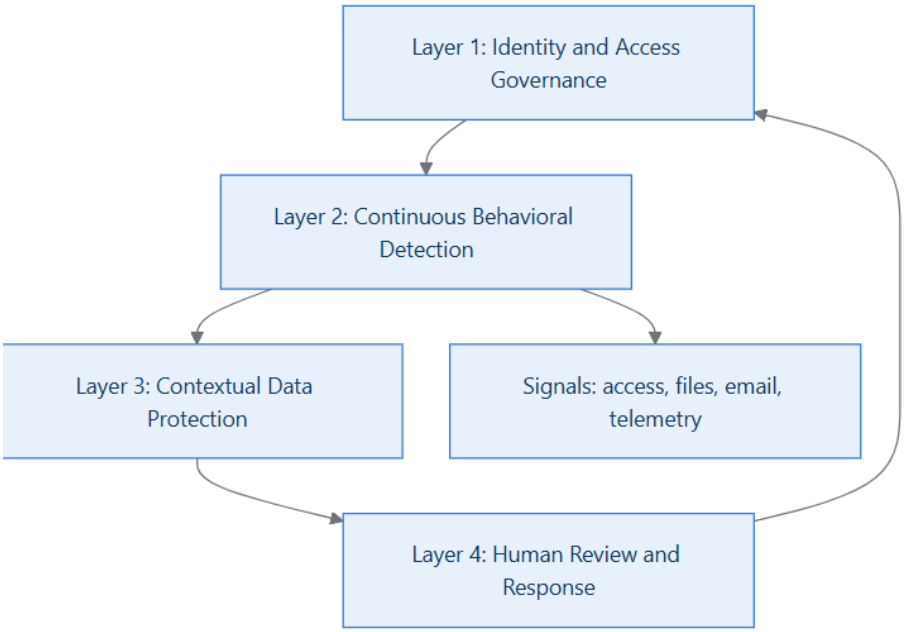

## 8. Challenges and Future Directions

Despite the potential of the framework described above, several challenges continue to limit insider threat detection. One of the biggest is the availability and quality of data. Insider threat datasets tend to be highly imbalanced, confirmed incidents are relatively rare, and employee behavior changes over time (Le et al., 2020). Although techniques such as oversampling can help address class imbalance (Al-Shehari & Alsowail, 2021), the lack of reliable ground truth remains a difficult problem.

Another challenge is adversarial adaptation. Insiders who understand security controls can change document wording to avoid keyword detection (Montano et al.,) or encrypt data before exfiltrating it to avoid content inspection (Ayobami et al.,). This means detection systems need to focus more heavily on signals that are harder to manipulate, such as metadata, endpoint telemetry, and broader behavioral patterns.

False positives are another concern. If security systems generate too many alerts that turn out to be harmless, analysts may begin to ignore or distrust the system. Static rule-based approaches can contribute to this problem by generating large numbers of false positives while still missing legitimate threats (Thapaliya & Gurung,).

Privacy and ethics also require careful consideration. Monitoring employee behavior can raise legitimate privacy concerns, particularly when organizations analyze sensitive communications or encrypted information. Future DLP and insider threat systems therefore need to become more intelligent without simply increasing the level of employee surveillance (Thapaliya & Gurung,).

Future research should focus on several related areas. These include privacy-preserving behavioral analytics, adaptive learning systems that can account for changes in user behavior (Le et al., 2020), closer integration between identity threat detection and automated response (Subramanyam,), and standardized datasets that allow researchers to make more meaningful comparisons between detection methods. This last issue is especially important because current studies often rely on different datasets, including synthetic data, which makes direct comparison difficult (Qawasmeh & AlQahtani,)(Al-Shehari & Alsowail, 2021).

## 9. Conclusion

Insider threats are different from many traditional cybersecurity threats because they take advantage of legitimate authorization rather than attempting to bypass it. Employees and other insiders can cause significant damage because they already have access to organizational systems and are generally trusted by the organization (Renaud et al.,). Malicious insiders therefore remain one of the most serious threats facing companies and government agencies (Le et al., 2020).

Traditional security controls such as firewalls, antivirus software, and static rule-based DLP were largely designed around external threats. As a result, they often lack the flexibility and behavioral awareness needed to identify insider activity in real time (Qawasmeh & AlQahtani,)(Thapaliya & Gurung,).

This paper shows that effective insider risk management depends on three connected capabilities. First, organizations need a clear understanding of different insider types and motivations because a premeditated data thief, a disgruntled employee, and a negligent employee require very different responses (Homoliak et al., 2019)(Weber et al., 2020)(Renaud et al.,). Second, organizations need continuous behavioral detection that uses individualized baselines, peer comparisons, and machine learning rather than relying entirely on static rules. Research has reported detection of up to 85% of malicious insiders at a 0.78% false positive rate, with some malicious behavior identified within 14 minutes of the first action (Le et al., 2020). Third, prevention needs to be built into the architecture itself through Zero Trust, just-in-time privileged access, and context-aware data protection. JIT access, for example, has been reported to reduce the privilege-attack window by more than 90% (Subramanyam,).

There are still important limitations, including class imbalance, limited ground truth, behavioral drift, evasion techniques, false positives, and legitimate employee privacy concerns (Le et al., 2020)(Montano et al.,)(Thapaliya & Gurung, 2025). Even so, the overall direction is clear. Insider risk should not be treated only as a detection problem or a compliance requirement. It is better understood as an ongoing security management challenge that requires identity governance, behavioral analytics, contextual data protection, and human judgment to work together as one adaptive system.

## References

1. Gheyas, I. A., & Abdallah, A. E. (2016). Detection and prediction of insider threats to cyber security: A systematic literature review and meta-analysis. Big Data Analytics, 1, Article 6. https://doi.org/10.1186/s41044-016-0006-0
2. Liu, L., De Vel, O., Han, Q.-L., Zhang, J., & Xiang, Y. (2018). Detecting and preventing cyber insider threats: A survey. IEEE Communications Surveys & Tutorials. https://doi.org/10.1109/COMST.2018.2800740
3. Ho, S. M., & Benbasat, I. (2018). Trustworthiness attribution: Inquiry into insider threat detection. Journal of the Association for Information Science and Technology, 69(2), 271–280. https://doi.org/10.1002/asi.23938

4. Moradpoor, N., Brown, M., & Russell, G. (2017). Insider threat detection using principal component analysis and self-organising map. In Proceedings of the 10th International Conference on Security of Information and Networks. ACM. https://doi.org/10.1145/3136825.3136859
5. Moriano, P., Pendleton, J., Rich, S., & Camp, L. J. (2017). Insider threat event detection in user-system interactions. In Proceedings of the 2017 ACM Workshop on Managing Insider Security Threats. ACM. https://doi.org/10.1145/3139923.3139928
6. Carvallo, P., Cavalli, A. R., & Kushik, N. (2017). Automatic derivation and validation of a cloud dataset for insider threat detection. In Proceedings of the 12th International Conference on Software Technologies (pp. 480–487). SciTePress. https://doi.org/10.5220/0006480904800487
7. Crawford, M., Peterson, G. L., Mills, R. F., & Grimaila, M. R. (2013). Insider threat detection using virtual machine introspection. In 2013 46th Hawaii International Conference on System Sciences (pp. 1821–1830). IEEE. https://doi.org/10.1109/HICSS.2013.278
8. Legg, P. A., Buckley, O., Goldsmith, M., & Creese, S. (2017). Automated insider threat detection system using user and role-based profile assessment. IEEE Systems Journal, 11(2), 503–512. https://doi.org/10.1109/JSYST.2015.2438442
9. Ko, L. L., Divakaran, D. M., Liau, Y. S., & Thing, V. L. L. (2017). Insider threat detection and its future directions. International Journal of Security and Networks, 12(3), 168–187. https://doi.org/10.1504/IJSN.2017.084391
10. Lo, O., Buchanan, W. J., Griffiths, P., & Macfarlane, R. (2018). Distance measurement methods for improved insider threat detection. Security and Communication Networks, 2018, Article 5906368. https://doi.org/10.1155/2018/5906368
11. Wall, D. S. (2013). Enemies within: Redefining the insider threat in organizational security policy. Security Journal, 26(2), 107–124. https://doi.org/10.1057/sj.2012.1
12. Silowash, G. (2013). Insider threat attributes and mitigation strategies (CMU/SEI-2013-TN-018). Software Engineering Institute, Carnegie Mellon University. https://doi.org/10.1184/R1/6574451.v1
13. CERT Insider Threat Team. (2013). Unintentional insider threats: A foundational study (CMU/SEI-2013-TN-022). Software Engineering Institute, Carnegie Mellon University. https://doi.org/10.1184/R1/6585575.v1
14. Lewellen, T., Silowash, G., & Costa, D. (2013). Insider threat control: Using plagiarism detection algorithms to prevent data exfiltration in near real time (CMU/SEI-2013-TN-008). Software Engineering Institute, Carnegie Mellon University. https://doi.org/10.1184/R1/6574484.v1
15. Yang, G., Ma, J., Yu, A., & Meng, D. (2016). Survey of insider threat detection. https://doi.org/10.19363/j.cnki.cn10-1380/tn.2016.03.003