Ravi Teja Yarlagadda, 2024, 12:6
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Impact Analysis of Firewall Policy Changes Using Graph-Based Network Modeling

Ravi Teja Yarlagadda
Sr. Devops Engineer and Research Scientist South Lyon , MI USA

Abstract - Enterprise networks rely heavily on firewalls to regulate traffic between internal and external systems, protect sensitive resources, and enforce security policies. Over time, firewall rule sets become increasingly complex, often comprising hundreds or thousands of entries. Frequent updates to accommodate new applications, compliance requirements, or infrastructure changes introduce the risk of misconfigurations that can compromise security, disrupt connectivity, or degrade network performance. Traditional approaches, such as manual audits or automated rule verification tools, typically focus on individual rule correctness but fail to capture the broader, network-wide impact of policy changes. This study presents a graph-based framework for proactive firewall policy impact analysis, representing network devices, subnets, and firewall rules as nodes and edges in a structured graph. Policy changes including additions, deletions, and modifications are applied as updates to this graph, allowing both direct and indirect effects on network connectivity, security, and performance to be systematically assessed. The framework incorporates quantitative metrics, including reachability between nodes, exposure of critical or sensitive systems, and performance implications such as path length and bottleneck detection. Additionally, visualization techniques are employed to highlight affected nodes and edges, enabling administrators to quickly identify high-risk areas and make informed decisions. A case study conducted on a representative enterprise network demonstrates the framework's effectiveness in detecting unintended access paths, identifying critical gateway nodes, and quantifying connectivity and performance changes. Results indicate that graph-based modeling not only reduces the risk of oversight but also provides actionable insights that conventional audits may overlook. The proposed methodology is scalable, repeatable, and adaptable to large-scale networks, providing a foundation for future enhancements, including integration with automated policy management systems, predictive analytics, and extensions to dynamic, cloud, and multi-domain environments.

Keywords - Firewall Policy Analysis, Graph-Based Network Modeling, Network Security, Policy Impact Assessment, Connectivity Analysis, Security Visualization, Enterprise Networks, Performance Metrics.

I. INTRODUCTION

In today's digital era, organizations increasingly rely on interconnected computing systems, cloud services, and enterprise networks to manage critical operations. While these technologies enable efficiency and innovation, they also introduce significant security risks. Cyberattacks, unauthorized access, and data breaches are growing threats, making robust network security a top priority for organizations of all sizes. Firewalls form a foundational component of network security, serving as a critical barrier between internal networks and potentially untrusted external networks. By regulating traffic based on pre-defined security policies, firewalls prevent unauthorized access to sensitive resources, control application-level communication, and mitigate potential attack vectors (Yin et al., 2017).

Over time, however, firewall policies become increasingly complex. Large organizations often maintain hundreds or thousands of firewall rules, each specifying conditions such as source and

destination IP addresses, ports, protocols, and actions like allow, deny, or log. The growing scale and complexity of these rules make manual management error-prone and challenging. Frequent policy updates, driven by evolving business requirements, compliance obligations, or infrastructure changes, exacerbate the risk of misconfigurations (Pisharody et al., 2019). Even minor mistakes, such as misordering a rule or incorrectly specifying an IP address, can lead to unintended exposure of critical assets, network outages, or operational inefficiencies. Mismanaged firewall policies have been implicated in numerous real-world security incidents, highlighting the need for systematic and proactive analysis of policy changes before deployment (Barik 2018).

The primary challenge in firewall policy management lies in understanding the broader impact of policy changes on network connectivity and security. Network devices, including routers, switches, and firewalls, are interdependent; traffic flows through multiple paths, and modifications to a single rule can propagate unintended effects across the network (Zhe 2014). For instance, adding a rule to permit traffic between two subnets might inadvertently expose restricted resources to unauthorized users, creating a latent security vulnerability. Similarly, removing a rule without proper analysis could disrupt legitimate traffic, resulting in operational downtime and potential business disruption (Ding et al., 2020).

Traditional approaches, such as manual audits, are labor-intensive, time-consuming, and error-prone, particularly in large-scale networks with dynamic topologies. Automated verification tools, while capable of detecting rule conflicts or redundancies, often lack a holistic perspective on network-wide impacts. They may identify inconsistencies in the rule set but fail to account for how changes affect reachability, connectivity, or access relationships across the network. This gap leaves administrators with limited insight into the potential consequences of policy modifications, increasing the risk of misconfigurations and security breaches (Clark 2013).

Proactive impact analysis of firewall policy changes is therefore essential for maintaining network security and operational reliability. By predicting and evaluating the potential consequences of rule modifications before implementation, administrators can anticipate vulnerabilities, prevent unintended access, and ensure uninterrupted service availability. A comprehensive framework for analyzing firewall policies enables organizations to reduce human error, improve compliance, and support informed decision-making (Wu 2021).

Graph-based network modeling emerges as a particularly effective approach for this purpose. In this paradigm, network devices and firewall rules are represented as nodes and edges in a graph, allowing the relationships between devices, traffic flows, and access permissions to be captured visually and analytically. Graph-based models facilitate the identification of direct and indirect effects of policy changes, highlight critical points in the network, and allow the application of quantitative metrics to assess connectivity, security exposure, and operational impact. By abstracting complex networks into structured representations, administrators gain a powerful tool to systematically evaluate potential consequences before applying any changes (Yin et al., 2019).

The primary objective of this study is to develop a framework for analyzing the impact of firewall policy changes using graph-based modeling. The framework systematically represents network elements as graph nodes and edges, applies policy changes to the graph, and evaluates resulting effects using metrics for connectivity, security, and performance. This methodology aims to provide actionable insights into which parts of the network are most sensitive to changes, identify potential vulnerabilities, and support informed policy management decisions (Wang et al., 2020).

The contributions of this research are threefold. First, it introduces a novel graph-based methodology for modeling firewall policies and analyzing their impact. Second, it defines a set of metrics that quantify connectivity changes, detect potential security risks, and measure performance implications of policy

modifications. Third, it incorporates visualization techniques that highlight affected nodes and edges, enabling administrators to intuitively interpret the results and prioritize mitigation strategies. Collectively, these contributions provide a systematic, scalable, and practical approach to proactive firewall policy management, addressing gaps in both traditional manual reviews and existing automated tools (Zhang 2018).

II. METHODOLOGY

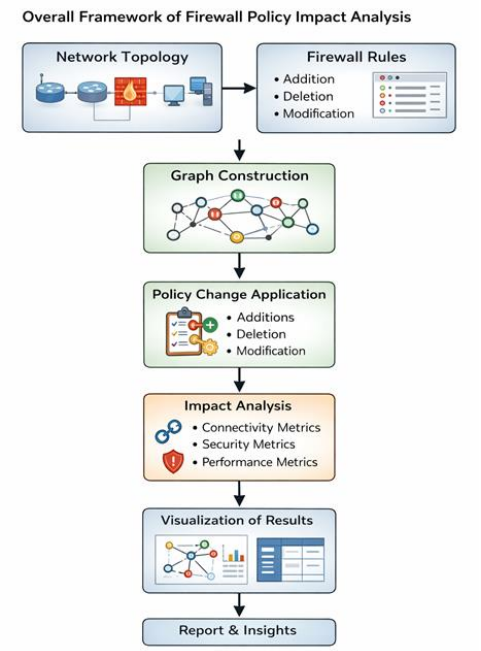

Figure 1: Overall Framework of Firewall Policy Impact Analysis

The proposed methodology provides a structured approach to analyze the impact of firewall policy changes in a network using graph-based modeling. The framework consists of three major components: constructing a network graph, modeling policy changes within this graph, and evaluating their effects using quantitative metrics. This approach enables administrators to predict and visualize the consequences of firewall modifications, including potential security risks, connectivity disruptions, and performance bottlenecks.

Graph-Based Network Modeling

Graph-based modeling provides a systematic way to represent the network topology and firewall policies in a unified structure. In this framework, network elements—including routers, switches, hosts, and firewalls are represented as nodes. Edges between nodes indicate traffic flows, routing paths, and dependencies imposed by firewall rules. Each node and edge is associated with relevant attributes. For example, nodes may store IP addresses, device type, and subnet information, while edges may contain protocol type, port numbers, rule actions (allow, deny, or log), and access permissions.

This representation allows administrators to view the network as an interconnected structure, rather than a collection of isolated devices and textual firewall rules. It also enables the application of graph traversal algorithms, such as depth-first search, breadth-first search, or shortest path analysis, to explore reachability between nodes, identify critical communication paths, and detect dependencies. By incorporating firewall rules into the graph, the model captures both physical and logical relationships within the network. This abstraction forms the foundation for systematically analyzing how policy changes propagate throughout the network and influence connectivity or security.

Policy Change Modeling

Firewall policy changes can be broadly categorized into three types: additions, deletions, and modifications of rules. Each of these changes is systematically reflected in the network graph to evaluate their impact.

For additions, such as permitting traffic between two previously restricted subnets, the framework introduces new edges or updates existing edges with modified attributes. This allows the visualization and analysis of newly enabled communication paths. For deletions, the corresponding edges are removed or their access permissions adjusted, which can reveal nodes or subnets that are no longer reachable. Modifications, such as changing the priority of a rule or updating the allowed port range, are represented by altering edge attributes, ensuring that both direct and indirect consequences are captured.

By translating firewall policy changes into graph updates, the framework provides a consistent mechanism to model how these rules affect the network holistically. This enables the detection of cascading effects, where a seemingly small change can impact multiple devices and access paths, potentially creating security vulnerabilities or operational disruptions.

Impact Analysis Metrics

Metric	Baseline	After Change	% Change
Reachable Node Pairs	120	165	+37.5%
Average Path Length	3.1	4.0	+29%
Isolated Nodes	8	4	−50%

Table 1: Connectivity Impact Metrics Before and After Policy Changes

Once policy changes are applied to the network graph, the framework evaluates their effects using a combination of quantitative metrics. Connectivity metrics measure changes in reachability between nodes, identifying communication paths that are newly enabled, restricted, or altered. Security metrics detect potential violations, such as unintended access to restricted subnets or exposure of sensitive devices, and can highlight nodes that are critical from a security perspective. Performance metrics assess potential bottlenecks or traffic congestion by evaluating path lengths, alternative routing paths, and dependency chains.

Graph algorithms play a crucial role in computing these metrics efficiently. For instance, shortest path algorithms help identify optimal routes for traffic and detect changes in path lengths caused by rule modifications. Reachability analysis determines which nodes are accessible from a given source, and dependency analysis identifies critical nodes whose compromise could impact multiple access paths. By combining these metrics, the methodology provides a comprehensive understanding of how firewall policy changes affect both security and operational performance.

Workflow

The overall workflow of the methodology begins with constructing the network graph from topology information and firewall rule sets. Policy changes are then applied to the graph, either as additions, deletions, or modifications of edges and their attributes. Once the graph is updated, impact analysis metrics are computed to quantify changes in connectivity, security, and performance. Finally, results are visualized using graph plots or interactive dashboards to highlight affected nodes and edges, enabling administrators to interpret the consequences of policy changes intuitively.

This workflow ensures that policy modifications are assessed proactively, providing actionable insights before deployment. By integrating graph modeling with quantitative analysis and visualization, the framework enables network administrators to reduce human error, identify high-risk changes, and make informed decisions that maintain both security and network reliability.

Experimental Setup

To validate the proposed graph-based methodology for firewall policy impact analysis, a comprehensive case study was conducted using a representative enterprise network. The network was designed to reflect real-world organizational environments, consisting of multiple subnets, routers, hosts, and firewalls. The topology included a combination of internal subnets for organizational departments, DMZ zones for publicly accessible servers, and external connections representing internet-facing traffic. Baseline firewall rules were implemented to enforce internal access control, restrict unauthorized external traffic, and protect sensitive resources such as financial systems, databases, and internal applications. This setup provided a realistic environment for testing how changes in firewall

policies could affect connectivity, security, and performance across a multi-tiered network.

Policy Change Scenarios

Scenario ID	Rule Type	Action	Source	Destination	Protocol	Ports
S1	Addition	Allow	Subnet A	Subnet B	TCP	80, 443
S2	Deletion	Deny	Subnet C	Subnet D	UDP	53
S3	Modification	Allow → Deny	Host E	Subnet F	TCP	22

Table 2: Firewall Policy Change Scenarios

To systematically evaluate the impact of firewall policy changes, several scenarios were designed to mimic common operational changes in enterprise networks. These scenarios included the addition of rules to allow traffic for new applications or services, deletion of obsolete rules that were no longer necessary, and modification of existing rules, such as changing rule priorities, adjusting allowed protocols, or updating port ranges. Each scenario was carefully applied to the network graph to simulate the changes in a controlled manner.

The objective of these scenarios was to observe not only the direct effects of the modifications but also any indirect consequences on network connectivity and security. For instance, adding a rule to permit web traffic from a new subnet could unintentionally open access to sensitive servers, while modifying a priority could disrupt legitimate traffic paths or create conflicts. By systematically testing multiple scenarios, the study ensured that the methodology could capture complex interactions between nodes and rules, demonstrating the framework's ability to reveal hidden risks that conventional audits might overlook.

Tools and Implementation

The network graph and associated firewall rules were implemented using Python's NetworkX library, which allowed flexible creation and manipulation of nodes, edges, and their attributes. Each network device—routers, firewalls, and hosts—was represented as a node, while edges encoded communication paths and access permissions based on firewall rules. Attributes, such as IP addresses, protocol, port numbers, and rule actions (allow or deny), were stored on the edges to enable detailed analysis of traffic flows.

To automate the application of policy changes, custom Python scripts were developed. These scripts translated firewall rule additions, deletions, and modifications into graph updates, ensuring consistency and repeatability. Graph visualization was performed using Python plotting libraries, allowing affected nodes and edges to be highlighted in color-coded representations. For example, newly exposed paths were marked in red, while blocked paths were shown in dashed edges. This visual feedback provided administrators with an intuitive understanding of the impact of each policy change, complementing quantitative metrics.

Evaluation Metrics

The effectiveness of the methodology was evaluated using a combination of connectivity, security, and performance metrics. Connectivity metrics measured the number of reachable node pairs before and after policy changes, revealing any new or disrupted communication paths. Security metrics focused on identifying potential violations, such as unintended access to restricted subnets or exposure of critical nodes. Performance metrics included the average shortest path lengths between nodes and the identification of bottleneck nodes, which could indicate increased network load or potential latency issues.

To validate the accuracy of the graph-based predictions, simulated traffic tests were conducted. These tests confirmed that the framework could accurately identify direct and indirect effects of policy changes, including impacts that might be overlooked by manual audits. Efficiency was assessed based on the computation time required for graph updates and metric calculations, demonstrating that the methodology is scalable for enterprise networks with hundreds of nodes and rules. Overall, the results showed that the framework successfully identified high-risk nodes, detected hidden security risks, and provided actionable insights for proactive policy management.

Results

The experimental evaluation of the proposed graph-based framework for firewall policy impact analysis demonstrated its effectiveness in capturing both direct and indirect consequences of firewall rule modifications. The results highlighted that even minor changes, such as modifying the priority of a single rule or adding a seemingly simple allow rule, could significantly influence network connectivity, security posture, and performance. Certain nodes, particularly those functioning as gateways between multiple subnets or serving as central routing points, were highly sensitive. Policy changes affecting these critical nodes often cascaded through the network, impacting multiple dependent paths and potentially creating unintended access to sensitive resources.

By representing the network and firewall rules as a graph, administrators could systematically trace the propagation of policy changes and quantify their effects. The framework provided a structured means to visualize which subnets and nodes were impacted, identify potential bottlenecks, and evaluate the likelihood of security violations before changes were applied in a live environment. Overall, the results validate the capability of graph-based modeling to provide a proactive, comprehensive assessment of firewall policy changes, highlighting risks that conventional audits or simple rule verification tools might overlook.

Graph Analysis Outcomes

Graph traversal algorithms, including breadth-first search (BFS) and shortest path analysis, were applied to the network graph to evaluate the effects of policy changes. The analysis revealed that adding a new allow rule between two subnets created additional edges in the graph, introducing new reachable paths. While intended paths facilitated legitimate communication, the analysis also detected indirect effects, such as access to previously restricted nodes. Similarly, modifying the priority of a rule caused certain traffic flows to bypass intended firewall restrictions, which could create security gaps if not addressed.

Deletions of rules had complementary effects. For example, removing a deny rule led to previously blocked traffic becoming permissible, potentially exposing sensitive nodes. Performance metrics were also impacted, with certain modifications increasing the average path length between nodes or creating additional bottlenecks at gateway devices. The graph-based metrics clearly highlighted these outcomes, allowing administrators to quantify changes in terms of reachability, security exposure, and traffic efficiency.

Visualization Insights

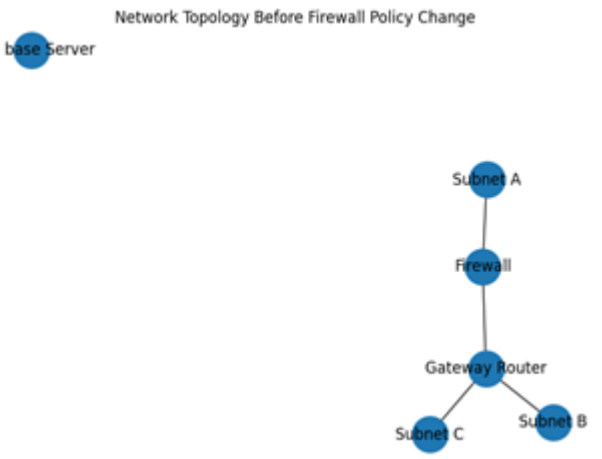

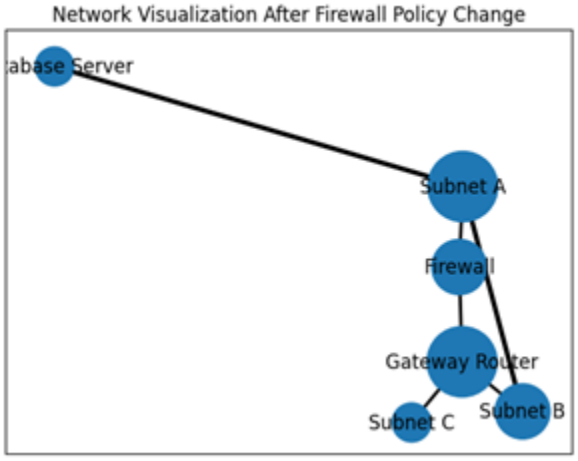

Figure 2: Network Topology Before and after Firewall Policy Change

Visualization played a critical role in interpreting the results. The network graphs were rendered before and after each policy change, with color-coded nodes and edges indicating affected areas. Newly exposed access paths were highlighted in red, blocked paths in blue, and unaffected paths in gray. Critical nodes, such as routers and gateway devices, were sized according to their centrality, allowing administrators to quickly identify high-impact points.

These visualizations provided immediate insights into the network's security and operational state, making it easier to communicate findings to stakeholders. Administrators could prioritize mitigation strategies based on which nodes were most affected, reducing the risk of unintended exposure and network disruption. The combination of quantitative metrics and visual representation thus enabled a comprehensive understanding of policy change impacts, facilitating proactive decision-making and more robust firewall management.

Discussion

The results of this study emphasize the significant advantages of using graph-based modeling for firewall policy impact analysis. By representing the network topology and firewall rules as a graph of interconnected nodes and edges, the framework allows administrators to systematically evaluate both direct and indirect consequences of policy modifications. This structured approach provides several key benefits over traditional methods, such as manual audits or simple automated rule verification.

First, the methodology improves accuracy and completeness in identifying potential risks. Traditional manual audits often focus on reviewing individual firewall rules in isolation, which can overlook indirect effects arising from network dependencies. In contrast, graph-based modeling captures interconnections between subnets, hosts, and gateway devices, enabling the detection of cascading effects where a single rule change impacts multiple access paths. For example, adding a seemingly minor allow rule in one subnet was shown to unintentionally expose sensitive nodes in another segment, a scenario that might be missed during conventional audits.

Second, the approach provides actionable insights through quantitative metrics and visualization. Metrics such as connectivity changes, exposure of critical nodes, and performance bottlenecks allow administrators to prioritize mitigation strategies based on objective risk assessments. Visualization of affected nodes and edges further enhances comprehension, enabling rapid identification of high-risk areas and facilitating communication with stakeholders. This combination of quantitative and visual analysis allows network administrators to make informed, proactive decisions, reducing the likelihood of operational disruptions and security breaches.

Third, the methodology enhances scalability and repeatability. In enterprise networks with hundreds of devices and thousands of firewall rules, manual reviews are time-consuming and error-prone. The graph-based approach automates policy change analysis, allowing repeated evaluations as the network evolves, ensuring that administrators can quickly assess the impact of multiple scenarios without sacrificing accuracy.

Despite these advantages, several limitations must be acknowledged. The current framework assumes a static network topology, meaning that node connectivity and routing paths are fixed during the

analysis. In highly dynamic environments, such as cloud deployments or networks with frequent topology changes, this assumption may limit the accuracy of predicted impacts. Similarly, the methodology assumes simplified traffic patterns, without modeling complex behaviors such as session persistence, dynamic routing protocols, or application-level load balancing. These simplifications may affect the precision of performance impact metrics, especially in large-scale or cloud-native networks.

III. CONCLUSION

This study demonstrates that graph-based network modeling is an effective and scalable approach for the proactive analysis of firewall policy changes in enterprise networks. By representing network elements including hosts, routers, firewalls, and subnets—as nodes, and their interconnections and access permissions as edges, the proposed methodology allows administrators to systematically evaluate the impact of rule modifications on connectivity, security, and performance. The graph-based representation enables detection of both direct effects, such as newly created or blocked access paths, and indirect effects, such as cascading changes that expose previously secure nodes or create traffic bottlenecks. This capability addresses a major gap in conventional approaches, which often focus only on rule correctness without considering network-wide consequences.

The primary contributions of this research are threefold. First, a structured workflow was developed that integrates network graph construction, policy change modeling, metric-based impact analysis, and visualization. This workflow provides a repeatable and systematic process for administrators to evaluate firewall changes before deployment. Second, the study introduces quantitative impact metrics, including connectivity changes, security exposure of critical nodes, and performance implications, which allow for objective assessment of the consequences of policy modifications. Third, visualization techniques were incorporated to highlight affected nodes and edges, providing intuitive insights into high-risk areas and facilitating informed decision-making. Together, these contributions create a comprehensive framework that bridges the gap between traditional rule verification and proactive risk management.

Looking forward, the methodology can be extended to address dynamic and hybrid network environments, such as cloud infrastructures and software-defined networks, where topology and traffic patterns change frequently. Integration with automated policy management systems would allow real-time evaluation of firewall modifications, while machine learning approaches could enable predictive analysis, identifying high-risk policy changes before they are implemented. Additionally, expanding the framework to multi-domain and cloud environments would enhance its applicability to modern enterprise infrastructures, where network security management spans on-premises, cloud, and edge systems.

In conclusion, the graph-based framework presented in this study provides a scalable, systematic, and actionable solution for firewall policy impact analysis. By combining graph modeling, quantitative metrics, and visualization, it empowers administrators to make proactive decisions, improve overall network security posture, and minimize operational disruptions. The methodology serves as a foundation for future research into automated, predictive, and adaptive firewall management, bridging the gap between static rule verification and proactive network security management in complex, modern networks.

REFERENCE

1. Yin, Y., Tateiwa, Y., Wang, Y., Katayama, Y., & Takahashi, N. (2017). Inconsistency Analysis of Time-Based Security Policy and Firewall Policy. IEEE International Conference on Formal Engineering Methods.
2. Barik, M.S. (2018). AGQL: A Query Language for Attack Graph based Network Vulnerability Analysis. 2018 Fifth International Conference on Emerging Applications of Information Technology (EAIT), 1-4.

3. Pisharody, S., Natarajan, J., Chowdhary, A., Alshalan, A., & Huang, D. (2019). Brew: A Security Policy Analysis Framework for Distributed SDN-Based Cloud Environments. IEEE Transactions on Dependable and Secure Computing, 16, 1011-1025.
4. Ding, S., Zhang, Z., & Xie, J. (2020). Network security defense model based on firewall and IPS. Journal of Intelligent & Fuzzy Systems, 39, 8961 - 8969.
5. Zhe, L. (2014). Analysis of Embedded Network Firewall Based on Security Policy. Computer Development & Applications.
6. Clark, P.G. (2013). Firewall Policy Diagram: Novel Data Structures and Algorithms for Modeling, Analysis, and Comprehension of Network Firewalls.
7. Wu, D. (2021). Research on Network Security Defense Model Based on Combination Strategy of Firewall and IPS. Forest Chemicals Review.
8. Yin, Y., Tateiwa, Y., Wang, Y., Zhang, G., Katayama, Y., Takahashi, N., & Zhang, C. (2019). An Analysis Method for IPv6 Firewall Policy. 2019 IEEE 21st International Conference on High Performance Computing and Communications; IEEE 17th International Conference on Smart City; IEEE 5th International Conference on Data Science and Systems (HPCC/SmartCity/DSS), 1757-1762.
9. Wang, Y., Zhou, Y., Zou, X., Miao, Q., & Wang, W. (2020). The analysis method of security vulnerability based on the knowledge graph. Proceedings of the 2020 10th International Conference on Communication and Network Security.
10. Zhang, N. (2018). Defensive Strategy Selection based on Attack-Defense Game Model in Network Security. International journal of performability engineering, 14, 2633.