Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 3, 2026

A Study on Firewall Effectiveness Against Modern Cyber Threats: A Critical Review of Technologies, Failure Modes, and Evaluation Approaches

[1]Ahmad Saim, [2]Faisal Mehmood

[1]BS Cyber Security, Riphah International University Faisalabad Campus
[2]BS Cyber Security, BS Cyber Security, Riphah International University Faisalabad Campus

[1]ahmadsaimmm123321@gmail.com, [2]jamfaisal146@gmail.com

DOI: https://doi.org/10.5281/zenodo.23480289

Keywords
Firewall; Next-Generation Firewall; Network Security; Zero Trust; Intrusion Prevention; TLS Inspection; Web Application Firewall; Cyber Threats; Configuration Errors

Article History
Received on 20 Jan 2026
Accepted on 15 Mar 2026
Published on 27 Mar 2026

Copyright @Author

Corresponding Author: *

Ahmad Saim
BS Cyber Security, Riphah International University Faisalabad Campus Email: ahmadsaimmm123321@gmail.com

Abstract

Firewalls have been a basic network security control for more than three decades, yet the environments they protect have changed substantially. Cloud adoption, remote work, pervasive encryption, and API-driven services have weakened the fixed network perimeter on which early firewall designs relied. This critical review examines how effectively contemporary firewall technologies, including packet filters, stateful inspection, application-layer gateways, next-generation firewalls (NGFWs), web application firewalls (WAFs), host-based and distributed firewalls, and cloud-delivered firewall services, counter present-day threats such as ransomware, web and application-layer attacks, distributed denial-of-service (DDoS) floods, encrypted command-and-control traffic, credential-based intrusions, and supply-chain compromise. Drawing on standards, peer-reviewed research, government advisories, and threat-intelligence frameworks, the review maps firewall capabilities to threat categories, synthesizes documented failure modes (configuration errors, rule-set decay, inspection evasion, encryption blind spots, and attacks against firewall appliances themselves), and discusses how firewalls fit within zero trust and defense-in-depth architectures. A seven-dimension framework for outcome-based evaluation and a four-level maturity model are proposed, together with practical recommendations and a research agenda. The review concludes that firewalls remain necessary but insufficient: their protective value depends less on product category than on policy quality, patch discipline, visibility into encrypted traffic, and integration with identity, endpoint, and monitoring controls.

INTRODUCTION

Background and motivation

Firewalls are among the oldest and most widely deployed network security controls. In its most general sense, a firewall is a device or software component that enforces an access-control policy on traffic passing between network segments, allowing or blocking communication according to defined criteria (Scarfone & Hoffman, 2009). The idea took shape in the late 1980s and early 1990s, when organizations began to connect internal networks to the public Internet and researchers were documenting weaknesses in the underlying TCP/IP protocols (Bellovin, 1989). Early practitioners described the firewall as a single choke point through which all traffic between a trusted inside and an untrusted outside had to pass, allowing security teams to concentrate defensive effort in one place (Cheswick et al., 2003).

The threat model behind that design has weakened considerably. Employees now work from home networks and personal devices, applications run across several public cloud providers, business data sits in software-as-a-service platforms, and partners connect through application programming interfaces (APIs) rather than dedicated links. The "inside" of the network is no longer a single trusted place, and the "outside" is no longer the only source of danger. At the same time, most web traffic is encrypted, which limits what devices in the network path can observe (Rescorla, 2018), and attackers frequently prefer stolen credentials, legitimate administrative tools, and trusted software channels to noisy exploitation of open ports.

These shifts have prompted a lively debate. Vendor literature often presents modern firewalls as comprehensive protection platforms, while some commentators argue that the perimeter is obsolete and that firewalls are largely irrelevant in a zero trust world (Kindervag, 2010; Rose et al., 2020). Both positions oversimplify. Firewalls continue to block large volumes of unwanted traffic and provide the segmentation on which many other controls depend, yet their limits are well documented, and the appliances themselves have become prominent targets for skilled adversaries (Cisco Talos, 2024; CISA, n.d.).

1.2 Problem statement

The practical question facing security managers, students, and researchers is therefore not whether firewalls are useful, but how useful they are, against which threats, under what conditions, and how that usefulness can be measured. Existing literature addresses parts of this question, including configuration errors (Wool, 2004, 2010), policy anomalies (Al-Shaer & Hamed, 2004), inspection evasion (Handley et al., 2001), TLS interception (Durumeric et al., 2017), and DDoS defense (Zargar et al., 2013). These strands are rarely brought together in a single assessment that maps technology capabilities to current threat categories and connects them to zero trust practice. Vendor-neutral, outcome-oriented evidence also remains scarce. This review responds to that gap.

Objectives and research questions

The overall objective is to provide a critical, evidence-based assessment of firewall effectiveness against modern cyber threats. The review is organized around four research questions:

RQ1. How have firewall technologies evolved, and what protective functions does each generation provide?

RQ2. How effective are firewalls against the major categories of contemporary threats?

RQ3. Which documented limitations and failure modes reduce firewall effectiveness in practice?

RQ4. How should effectiveness be evaluated, and what role should firewalls play in zero trust and defense-in-depth architectures?

Contributions

The article makes three contributions. First, it presents a capability-to-threat analysis across

firewall generations, supported by a qualitative comparison matrix. Second, it consolidates evidence on failure modes, including configuration decay, inspection evasion, encryption blind spots, and the exposure of firewalls as attack targets. Third, it proposes a seven-dimension evaluation framework and a maturity model that emphasize measurable outcomes over feature lists, and it offers recommendations for organizations of different sizes together with an agenda for future research.

Organization of the paper

Section 2 explains the review approach and scope. Section 3 traces the evolution of firewall technologies. Section 4 characterizes the modern threat landscape. Section 5 analyzes effectiveness by threat category, and Section 6 examines how effectiveness has been measured in the literature. Section 7 discusses structural limitations and failure modes, Section 8 considers specialized environments, and Section 9 reviews emerging directions. Section 10 proposes the evaluation framework, Section 11 gives recommendations, Section 12 identifies research gaps, Section 13 notes the limitations of the review, and Section 14 concludes.

Review Approach and Scope

This article is a narrative critical review rather than a systematic review or meta-analysis. The heterogeneity of the available evidence, which ranges from standards documents and controlled experiments to configuration audits, vulnerability advisories, and threat-intelligence reports, makes statistical pooling inappropriate. A narrative approach allows different kinds of evidence to be weighed in context and allows the reviewer to explain why a particular source is or is not persuasive.

Sources were drawn from four groups: (a) standards and guidance issued by bodies such as the National Institute of Standards and Technology (NIST) and the Internet Engineering Task Force (IETF); (b) peer-reviewed research on network security, including studies of configuration quality, policy analysis, intrusion detection, encrypted-traffic analysis, and DDoS defense; (c) vulnerability databases, government advisories, and vendor-neutral threat reports; and (d) established textbooks and technical reports.

Works were included when they addressed the design, evaluation, configuration, or security of firewalls or closely related enforcement points; described threats relevant to the environments in which firewalls operate; or offered methodological insight into how network security controls can be measured. Marketing material and unverifiable performance claims were excluded. Vendor documentation was used only to describe product behavior or vulnerabilities, not to support efficacy claims.

The review concentrates on enterprise, campus, and cloud environments. Industrial control systems and consumer routers are discussed only briefly. Because threat statistics change quickly, the article emphasizes structural relationships rather than numerical prevalence figures, and readers seeking current numbers should consult the latest editions of annual breach and threat reports.

Conceptual Background and Evolution of Firewall Technologies

Purpose and core principles

A firewall performs several functions: it filters traffic against a rule set, records connection events, often translates network addresses, frequently terminates virtual private network (VPN) tunnels, and separates networks into zones of differing trust. Its security value rests on three principles that recur throughout the literature and in official guidance: default-deny (permit only what is explicitly required), least privilege (grant each source only the access it needs), and choke-point mediation (ensure that traffic between zones actually passes through an enforcement point) (Scarfone & Hoffman, 2009; Cheswick et al., 2003). A history of the

technology is given by Ingham and Forrest (2002); the following subsections summarize the main generations.

Packet-filtering firewalls

The earliest firewalls examined each packet independently against static rules based on source and destination addresses, ports, and protocols. Packet filters are fast, inexpensive, and easy to implement in routers, which explains their persistence. Their weakness is the absence of context: they cannot tell whether a packet belongs to a legitimate session, cannot reassemble fragments in any meaningful way, and cannot inspect payloads (Cheswick et al., 2003). Rules for return traffic often have to be written broadly, which opens more of the network to unsolicited connections than administrators intend.

Stateful inspection

Stateful firewalls maintain a table of active connections and permit return traffic only for sessions that were legitimately initiated from a permitted direction. This closed many spoofing and unsolicited-inbound gaps, reduced the number of rules required, and became the foundation of most subsequent products (Scarfone & Hoffman, 2009). Stateful tracking also supports simple resistance to some resource-exhaustion attacks, for example through connection limits and SYN-flood protections, although its connection tables are themselves finite resources that flooding attacks can exhaust.

Proxy and application-layer gateways

Proxy-based designs terminate the client's connection, evaluate the request at the application layer, and open a separate connection to the destination. This allows deep validation of protocol behavior, concealment of internal hosts, and strong authentication and logging. The costs are added latency, reduced throughput, and the need for protocol-specific proxies. Proxy architectures remain valuable for high-assurance zones and for outbound web filtering, and their ideas were absorbed into later products (Cheswick et al., 2003).

Next-generation firewalls

The term next-generation firewall (NGFW) was popularized by industry analysts in the late 2000s to describe products that combine stateful filtering with application identification, integrated intrusion prevention, user-identity awareness, URL and content filtering, and often sandboxing or threat-intelligence feeds. The essential change is that policy can be written in terms of applications and users rather than only ports and addresses, which matters because many applications now share a small number of ports. Intrusion prevention within NGFWs descends from earlier network intrusion detection systems such as Snort and Bro (now Zeek), which inspect traffic for known patterns or anomalous behavior (Roesch, 1999; Paxson, 1999). The consolidation of functions simplifies procurement and management, but it also concentrates risk in a single platform and increases the performance cost of full inspection.

Web application firewalls

Web application firewalls (WAFs) protect web applications and APIs by inspecting HTTP and HTTPS requests and responses for injection attempts, cross-site scripting, malformed input, and abusive request patterns. They may be deployed as appliances, reverse proxies, or cloud services, and community rule sets such as the OWASP Core Rule Set provide a baseline. WAFs address a different layer than network firewalls and therefore complement rather than replace them. They are also a common form of virtual patching, providing temporary protection while application code is corrected (OWASP Foundation, 2021).

Host-based, distributed, and microsegmentation approaches

Bellovin (1999) argued that enforcing policy only at the network boundary is fragile and proposed distributed firewalls in which a central policy is enforced at each host. This idea anticipated today's host firewalls, endpoint security suites, and microsegmentation platforms, which restrict east-west traffic between workloads. Software-defined networking has further enabled programmable, centrally managed enforcement inside data centers (Kim & Feamster, 2013). The trade-off is management complexity: policies must be distributed, verified, and kept consistent across many enforcement points, and the security of the endpoint on which the firewall runs becomes a dependency.

Cloud-native and firewall-as-a-service models

Cloud providers offer native controls such as security groups and network access control lists, and third-party vendors offer virtual appliances and firewall-as-a-service (FWaaS) platforms. The idea that network functions can be outsourced to the cloud predates the current market (Sherry et al., 2012). Delivering inspection from cloud points of presence allows consistent policy for remote users and branch offices without routing all traffic through a central data center. Such services are often bundled with secure web gateways and zero trust network access under the label secure access service edge (SASE). Table 1 summarizes the main technologies.

Table 1

Comparison of major firewall technologies

Technology	Inspection depth	Typical placement	Principal strengths	Principal limitations
Packet filter	Network and transport headers; stateless	Routers, basic gateways	Speed, simplicity, low cost	No connection context; no payload inspection; broad return-traffic rules
Stateful inspection	Headers plus connection state	Perimeter and data-center edge	Blocks unsolicited inbound traffic; efficient; well understood	Limited application awareness; cannot examine encrypted payloads
Proxy / application gateway	Full protocol at the application layer	Perimeter and high-assurance zones	Deep protocol validation; hides internal hosts	Latency; per-protocol support; limited scalability

Technology	Inspection depth	Typical placement	Principal strengths	Principal limitations
Next-generation firewall	Application, user, and content; integrated IPS	Perimeter and internal segmentation	Application- and user-based policy; consolidated controls	Cost and complexity; performance loss under full inspection; encryption blind spots
Web application firewall	HTTP(S) requests and API calls	In front of web applications and APIs	Blocks common injection and abuse patterns; virtual patching	Tuning burden; business-logic flaws; bypass through obfuscation
Host-based / distributed	Traffic at each endpoint	Servers, workstations, virtual machines	Fine-grained east-west control; travels with the workload	Management at scale; depends on endpoint integrity
Cloud-native / FWaaS	Varies from layer 3 to layer 7	Virtual networks and cloud edge	Elastic scale; consistent policy for remote users	Provider dependency; shared-responsibility ambiguity; latency and data-residency concerns

Note. Compiled by the authors from Scarfone and Hoffman (2009), Cheswick et al. (2003), Bellovin (1999), and Sherry et al. (2012).

The Modern Threat Landscape

Frameworks for describing adversary behavior

Two conceptual tools are widely used to describe attacks. The intrusion kill chain divides an operation into stages (reconnaissance, weaponization, delivery, exploitation, installation, command and control, and actions on objectives) and argues that defenders gain an advantage by disrupting any single stage (Hutchins et al., 2011). The MITRE ATT&CK knowledge base catalogs observed adversary tactics and techniques and provides a common vocabulary for mapping defensive coverage (MITRE, n.d.). Both are useful for firewall evaluation because they reveal which stages a firewall can plausibly influence, chiefly reconnaissance, delivery, exploitation of exposed services, command and control, and exfiltration, and which occur beyond its view, such as credential theft on an

endpoint or privilege escalation inside an operating system.

Ransomware and extortion

Ransomware operations typically follow a recognizable path: initial access through phishing, stolen or purchased credentials, or exploitation of internet-facing systems; discovery and lateral movement; privilege escalation; data theft; and finally encryption and extortion. Threat reports from European and industry sources have consistently placed ransomware among the most significant threats to organizations (ENISA, 2023; Verizon, 2024). Firewalls influence the early and late phases by limiting exposed services, blocking known malicious destinations, and restricting outbound data flows. If internal segmentation is enforced, they also influence the middle phase. They have little influence over the phishing message that delivers a malicious attachment or over a credential that an attacker has bought from a criminal marketplace.

Exploitation of edge infrastructure

A distinctive feature of recent years is that attackers increasingly target the very devices intended to protect the network. VPN concentrators, firewalls, and gateways sit at the boundary, are reachable from the Internet, hold privileged access, and often support only limited endpoint monitoring, which makes them attractive footholds. CISA's Known Exploited Vulnerabilities catalog repeatedly lists network edge products (CISA, n.d.), and annual breach investigations have reported growth in vulnerability exploitation as an initial-access route (Verizon, 2024). A compromised perimeter device can give an attacker traffic visibility, credentials, and a trusted position from which to reach internal systems.

Web and API attacks

Applications remain a favored target because they are exposed by design. The OWASP Top 10 documents recurring classes of weakness such as broken access control, injection, and security misconfiguration (OWASP Foundation, 2021). The Log4Shell vulnerability (NVD, CVE-2021-44228) illustrated how a flaw in a widely embedded library could allow remote code execution through crafted input and how quickly exploitation attempts can spread once a vulnerability is public. Firewalls and WAFs can shield exposed applications in the interim, but the durable fix is to patch the software.

Encrypted command-and-control and exfiltration

Transport Layer Security (TLS) 1.3 removes static RSA key exchange, mandates forward secrecy, and encrypts more of the handshake, including server certificates (Rescorla, 2018). These improvements protect users, but they also mean that passive monitoring devices cannot decrypt recorded traffic using a server's long-term key and can see less handshake metadata than before. Malware authors exploit the same protections. Command-and-control channels and data-exfiltration flows routinely use HTTPS or other encrypted protocols so that they blend with ordinary business traffic.

Distributed denial-of-service attacks

Mirkovic and Reiher (2004) classified DDoS attacks by their mechanism, and later work distinguishes volumetric floods that saturate bandwidth, protocol attacks that exhaust connection state, and application-layer floods that consume server resources (Zargar et al., 2013). The Mirai botnet showed how large numbers of poorly secured internet-connected devices can supply attack capacity on a massive scale (Antonakakis et al., 2017; Kolias et al., 2017). Threat reports continue to identify availability attacks as a persistent category (ENISA, 2023).

Credential-based and insider threats

Attacks that use valid credentials from valid networks are among the hardest to detect at the

network layer, because the traffic is authorized in every technical sense. The same is true of malicious or careless insiders who already hold legitimate access. Defending against them requires identity governance, multi-factor authentication, privileged-access management, endpoint telemetry, and behavioral analytics, with firewalls contributing mainly through segmentation and logging.

Supply-chain compromise

Supply-chain attacks insert malicious code into trusted software or update channels. The compromise of a widely used network-management product in 2020, which affected government and private-sector organizations, demonstrated that malicious updates delivered from trusted vendors are typically permitted by policy (CISA, 2020). Firewalls can limit the damage by restricting which internal systems may communicate with the Internet and by flagging unexpected destinations, but they cannot distinguish a legitimate update from a malicious one delivered over the same channel.

Effectiveness Analysis by Threat Category

Network-layer threats

Firewalls are most reliable at the network layer. Default-deny policies, stateful tracking, and segmentation prevent unsolicited connections, reduce the attack surface, and constrain the paths available to an intruder. These controls remain among the most cost-effective in the security toolkit and are foundational in official guidance (Scarfone & Hoffman, 2009). A recurrent pattern in incident analysis is the exposure of services that had no business reason to be reachable from the Internet; a well-maintained firewall policy that closes such exposure removes a whole class of attacks before any inspection engine is involved.

Exploit prevention and inspection evasion

Integrated intrusion prevention adds detection of known exploit patterns and can provide short-term protection for unpatched systems. Signature-based approaches are, however, inherently reactive: they detect what has been described. Research on inspection evasion has long shown that ambiguities in how network devices and end hosts reassemble fragments and streams can be exploited to hide attacks from an inspection device while still delivering them to the target (Ptacek & Newsham, 1998). Handley et al. (2001) proposed traffic normalization as a defense, in which the inspection device resolves ambiguities before forwarding traffic, and demonstrated the approach's feasibility while acknowledging its costs. Normalization reduces, but does not eliminate, the evasion problem, and intrusion prevention should be regarded as a compensating control rather than a substitute for patching.

Web and API attacks

WAFs can block many common injection and cross-site scripting attempts and can be updated rapidly when a new vulnerability class emerges. Their effectiveness depends on tuning. Overly strict rules break legitimate application functions and generate false positives, while permissive rules leave gaps, and payload obfuscation can bypass pattern-based rules. WAFs are also poorly suited to business-logic flaws, such as manipulating a discount workflow or accessing another user's record through a valid request, because the individual requests look legitimate. Secure development, code review, and application-level testing therefore remain essential.

Encrypted traffic and TLS inspection

TLS inspection, in which the firewall acts as an intermediary that decrypts, inspects, and re-encrypts traffic, restores visibility but brings costs: computational overhead, certificate-management complexity, privacy and regulatory concerns, and incompatibility with applications that use certificate pinning. Durumeric et al. (2017) measured HTTPS interception in the wild and found that

interception products frequently reduced connection security, in some cases introducing serious vulnerabilities. This finding does not argue against inspection, but it shows that poorly implemented interception can create risks larger than those it removes.

A complementary line of research investigates detecting malicious encrypted traffic without decryption. Anderson and McGrew (2016) and Anderson et al. (2018) showed that features such as TLS handshake parameters, packet sizes, and timing can help distinguish malicious from benign encrypted flows. Such approaches depend on representative training data and may degrade as traffic patterns and malware behavior change. In practice, organizations often combine selective decryption of higher-risk categories with metadata analysis, DNS controls, and endpoint telemetry.

Volumetric and application-layer DDoS

Enterprise firewalls at the network edge have limited capacity to absorb volumetric floods; once the upstream link is saturated, no on-premises device can help. Effective mitigation depends on upstream scrubbing services, content delivery networks, or provider-level filtering, combined with rate limiting and resilient application design (Zargar et al., 2013). Firewalls still contribute by dropping malformed traffic, limiting per-source connection rates, and protecting their own state tables, but they should not be expected to withstand large attacks alone.

Lateral movement and ransomware containment

Internal segmentation and microsegmentation address the lateral-movement phase of intrusions. When workstations cannot communicate with each other and servers accept only the flows that their functions require, an attacker who compromises one host has far fewer paths to exploit. This is directly relevant to ransomware, whose impact depends on the number of systems the operator can reach. The effectiveness of segmentation depends on how well policies reflect real application dependencies; overly permissive "any-any" rules within a segment defeat the purpose. Quantitative comparisons of segmentation strategies for ransomware containment are still scarce (see Section 12).

Machine-learning-based detection

Many products now advertise machine-learning detection. Sommer and Paxson (2010) explained why anomaly detection is harder in network security than in many other machine-learning applications: the cost of false positives is high, normal traffic changes continuously, attacks are rare and diverse, and representative labeled data are difficult to obtain. Evaluation datasets have themselves been criticized. McHugh (2000) analyzed weaknesses in the widely used DARPA intrusion-detection evaluations, and Tavallaee et al. (2009) identified problems in the KDD Cup 99 data set, while newer data sets such as CICIDS2017 were created to reflect more current traffic (Sharafaldin et al., 2018). Models tested on such data may not generalize to production networks, so vendor claims should be validated in the target environment.

Synthesis

Table 2 summarizes the assessment by threat category, and Figure 1 presents a qualitative comparison across technologies. Three patterns emerge. First, effectiveness is highest where the relevant characteristics are explicit in packet headers or connection behavior, as in scanning and unsolicited inbound access. Second, effectiveness is moderate where content can be inspected and patterns are known in advance, as in exploits against exposed services and common web attacks; here, evasion, signature lag, and encryption reduce reliability. Third, effectiveness is low where malicious activity is indistinguishable from authorized activity at the network level, as in credential misuse, insider abuse, and trusted-channel supply-chain

compromise, and where the attack exceeds the device's capacity, as in large volumetric DDoS.

Table 2
Firewall contribution, effectiveness, and limitations by threat category

Threat category	Firewall contribution	Effectiveness (qualitative)	Principal limitations	Complementary controls
Scanning and unsolicited inbound access	Default-deny and stateful filtering block most attempts	Strong	Only as good as rule discipline and minimization of exposed services	Asset inventory; attack-surface management
Known exploits against exposed services	IPS signatures; virtual patching	Moderate to strong	Signature lag; evasion; encrypted payloads	Timely patching; vulnerability management
Web and API attacks	WAF rules for injection and abuse	Moderate	Business-logic flaws; novel payloads; tuning burden	Secure development; application testing
Ransomware (network stages)	Segmentation; egress filtering; blocking known C2	Moderate	Phishing and endpoint execution occur outside firewall view	Endpoint detection; backups; identity controls
Encrypted C2 and exfiltration	Selective TLS inspection; metadata and DNS controls	Low to moderate	Privacy, performance, and pinning constraints	Endpoint telemetry; DNS security; behavioral analytics
Volumetric DDoS	Rate limiting; protection of state tables	Low at the enterprise edge	Link saturation cannot be resolved locally	Upstream scrubbing; CDN; provider filtering
Credential misuse and insider threats	Segmentation and logging	Low	Traffic appears legitimate	MFA; privileged-access management; user behavior analytics

Threat category	Firewall contribution	Effectiveness (qualitative)	Principal limitations	Complementary controls
Supply-chain compromise	Egress restriction; anomaly flagging	Low to moderate	Trusted channels are allowed by design	Software integrity verification; vendor risk management

Note. Qualitative synthesis by the authors from the sources cited in Sections 4 and 5. Ratings indicate general tendencies, not measured detection rates.

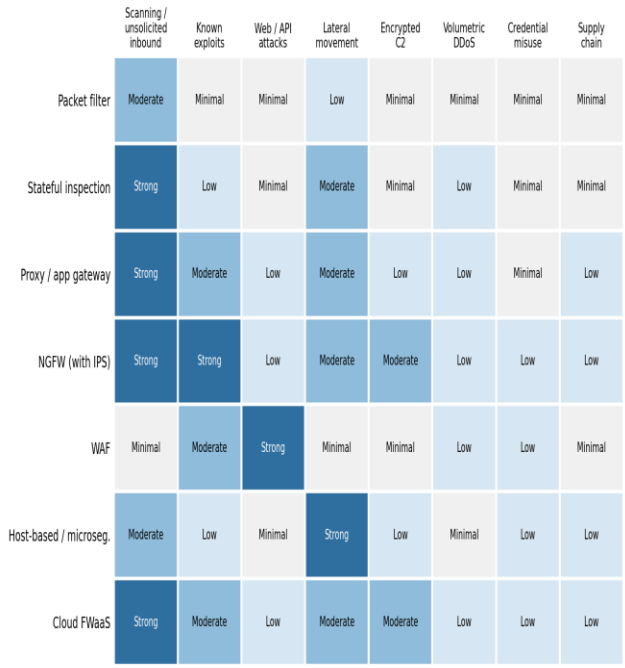


Figure 1

Qualitative comparison of firewall technologies against major threat categories. The ratings are an illustrative synthesis of the literature reviewed in this article and do not represent empirical measurements.

How Firewall Effectiveness Has Been Measured

Configuration-centric evaluation

Because a firewall enforces whatever policy it is given, a large body of research evaluates the quality of the policy rather than the device. Wool (2004) analyzed real-world rule sets and found configuration errors to be pervasive and correlated with rule-set complexity, and a follow-up study reported that errors remained common several years later and that complexity remained a strong predictor (Wool, 2010). Al-Shaer and Hamed (2004) formalized policy anomalies such as shadowing, redundancy, correlation, and generalization, and Hamed and Al-Shaer (2006) extended the analysis to conflicts across network security policies. Hu et al. (2012) proposed methods to detect and resolve such anomalies, while Gouda and Liu (2007) and Liu and Gouda (2008) proposed structured and diverse firewall design methods intended to reduce specification errors. These works treat effectiveness as a property of policy correctness and consistency.

Verification and reachability analysis

Formal and model-based approaches check whether the combined effect of routers, firewalls, and other devices matches the intended reachability policy. Al-Shaer et al. (2009) demonstrated an end-to-end verification approach that considers the interactions of network configurations and security policies. Such methods are powerful because they can detect errors that no single-device audit would reveal, but they require accurate models of the network and are less common in operational practice than manual audits.

Functional and performance testing

Functional tests assess whether a device blocks a defined set of attacks and how much throughput it sustains with various features enabled. Published throughput figures may be measured under favorable traffic mixes or with

some inspection features disabled, so independent replication under realistic conditions is valuable. Comparisons across products are further complicated by differences in default policies, signature sets, and tuning effort.

Threat-informed testing

Adversary-emulation and purple-team exercises use documented attacker techniques, for example those cataloged in MITRE ATT&CK, to test whether controls block or detect specific behaviors (MITRE, n.d.). Breach-and-attack-simulation tools automate parts of this process. The advantage is realism and traceability to real adversary behavior. The limitations are cost, dependence on the tester's skill, and the difficulty of generalizing from a small number of scenarios.

Limits of current metrics

Three problems recur. First, block rates depend heavily on the test set, so a high rate against a curated set of known attacks says little about resilience against novel ones. Second, the absence of alerts is not evidence of the absence of attacks, especially where encrypted or legitimate-looking traffic is involved. Third, outcome measures such as dwell time or incident cost are influenced by many controls at once, which makes it difficult to attribute improvement to the firewall. These difficulties motivate the multi-dimensional framework proposed in Section 10.

Structural Limitations and Failure Modes

Misconfiguration and policy decay

Misconfiguration is arguably the most consistent finding in the empirical literature. Rule sets grow as new applications, temporary exceptions, and emergency changes accumulate, while decommissioned systems leave orphaned rules behind. Redundant, shadowed, and overly permissive rules increase both the attack surface and the difficulty of review (Al-Shaer & Hamed, 2004; Wool, 2010). Changes made under operational pressure without peer review or testing are a common cause. In practice, a misconfigured advanced firewall may protect less than a carefully managed basic one, which is why evaluation should begin with policy quality rather than product features.

Firewalls as attack targets

Because firewalls and VPN gateways are internet-facing, privileged, and often poorly monitored internally, they attract sophisticated attackers. Table 3 lists selected examples that illustrate the pattern. In each case, exposure of a management or remote-access service on the device, combined with delays in patching or difficulty detecting compromise, played a role in the outcomes described in the advisories.

Table 3

Selected vulnerabilities and campaigns affecting perimeter security devices

Item	Affected product	Nature of the issue	Significance for firewall effectiveness
CVE-2022-42475	Fortinet FortiOS SSL-VPN	Heap-based buffer overflow allowing unauthenticated remote code execution; exploited in the wild	Shows how a remote-access service on a firewall can become the initial-access vector

Item	Affected product	Nature of the issue	Significance for firewall effectiveness
CVE-2024-3400	Palo Alto Networks PAN-OS GlobalProtect	Command injection allowing unauthenticated code execution with elevated privileges under specific configurations; exploited in the wild	Shows that an enabled feature on the firewall itself can create a critical exposure
ArcaneDoor campaign (2024)	Cisco ASA and other perimeter network devices	Espionage-focused campaign that compromised perimeter devices and maintained persistence	Shows sustained targeting of network infrastructure that often receives limited monitoring

Note. Based on NVD entries for CVE-2022-42475 and CVE-2024-3400 and on the Cisco Talos (2024) advisory. Readers should consult the current vendor advisories for patch status and mitigations.

These cases have a common lesson: a firewall is software running on an operating system, with the same potential for defects as any other exposed system. Management interfaces should not be reachable from the Internet, remote-access features should be enabled only where required and patched urgently, logs and integrity should be monitored, and organizations should plan for the possibility that the device itself is compromised.

Erosion of the perimeter

When users, data, and workloads sit outside the corporate boundary, a central chokepoint cannot observe or control most activity. Kindervag (2010) described the resulting "hard shell, soft center" weakness of perimeter-centric models and argued that trust should not be granted on the basis of network location. Google's BeyondCorp initiative described a practical migration from a privileged corporate network toward access decisions based on user and device state (Ward & Beyer, 2014). These developments do not make firewalls irrelevant, but they reduce the fraction of relevant traffic that any single firewall can see.

Operational overhead and alert fatigue

Feature-rich NGFWs generate large volumes of logs and alerts. Without adequate triage capacity, meaningful detections can be lost among noise, and staff may respond by disabling noisy signatures, thereby reducing protection. The skills required to tune intrusion prevention, manage certificates for TLS inspection, and maintain complex policies are scarce in many organizations, which magnifies the risk of error.

Performance and security trade-offs

Enabling deep inspection, TLS decryption, and sandboxing can significantly reduce throughput. Operators facing latency complaints may disable protections on high-volume segments or exempt large categories of traffic. Capacity planning that assumes full inspection from the outset, and periodic review of exceptions, are needed to prevent performance pressure from quietly eroding security.

Opacity and vendor dependence
Detection logic in commercial products is generally proprietary, which limits independent evaluation and makes it difficult for customers to know what a given signature or machine-learning model does and does not cover. Dependence on a single vendor for enforcement, threat intelligence, and management also concentrates risk: a defect in the platform, a faulty update, or a service outage can affect the whole environment.

Organizational and human factors
Technical controls operate within organizations. Unclear ownership of rules, "temporary" exceptions that become permanent, fragmented responsibility between network and security teams, and inadequate change management all contribute to policy decay. Studies of configuration quality suggest that process discipline, such as mandatory review, documentation of business justification for each rule, and expiry dates for exceptions, is as important as technical capability (Wool, 2010).

Firewalls in Specialized Environments

Cloud and hybrid environments
In cloud environments, network controls are largely software-defined, workloads are ephemeral, and responsibility is shared between provider and customer. Native security groups make it easy to define and change rules, but they also make it easy to create overly permissive rules, for example by exposing administrative ports to the whole Internet. Rule sprawl can reappear in a new form as the number of accounts, virtual networks, and services grows. Consistent policy across on-premises and cloud environments, infrastructure-as-code templates that embed secure defaults, and continuous configuration monitoring are the main mitigations.

Remote work and branch connectivity
Remote and mobile users often bypass the corporate firewall unless traffic is deliberately tunneled to it. Traditional VPN designs that route all traffic through a central site create performance and capacity problems and place a high-value VPN gateway on the Internet, as the incidents in Table 3 show. SASE and zero trust network access offerings address this by enforcing policy closer to the user and by granting access to specific applications rather than to whole networks, though they introduce dependency on the provider's service and require careful evaluation of what is inspected and where.

Internet of Things and operational technology
Many IoT and operational-technology devices cannot run endpoint agents, are rarely patched, and may speak legacy protocols without authentication. Network-level segmentation and firewalling are therefore particularly valuable, since they may be the only practical control. The Mirai episode showed the consequences of leaving such devices directly exposed (Antonakakis et al., 2017). Effective practice includes isolating these devices on dedicated segments, allowing only required flows, and applying protocol-aware inspection where available.

Resource-constrained organizations
Small and medium-sized enterprises, universities, and public-sector bodies in developing economies often face limited budgets and a shortage of specialized staff. For such organizations, the most valuable measures are generally the least glamorous: a default-deny policy, minimal exposure of services, prompt patching of internet-facing systems, multi-factor authentication, tested off-line backups, and log retention. Managed firewall services and well-maintained open-source firewall platforms can provide capable enforcement at modest cost, provided that

someone is clearly responsible for updates and rule review. Purchasing an advanced NGFW without the capacity to configure and monitor it may yield less protection than a simpler, well-run deployment.

Emerging Directions

Integration with zero trust architecture

NIST's zero trust architecture describes a model in which no request is trusted by virtue of network location; access decisions are made per session by a policy engine that considers identity, device posture, and context, and policy enforcement points are placed close to the resources they protect (Rose et al., 2020). In this model firewalls become one kind of enforcement point among several. Their strengths, including granular network policy, logging, and integration with identity, remain valuable, but they are no longer the primary defense. Gilman and Barth (2017) and the NIST Cybersecurity Framework 2.0 (NIST, 2024) similarly stress layered controls and continuous verification.

Microsegmentation and software-defined enforcement

Fine-grained east-west controls reduce the blast radius of compromised hosts and directly target the lateral-movement phase common to ransomware operations. Software-defined networking and host-based policy engines make it feasible to derive rules from observed application dependencies and to enforce them consistently (Kim & Feamster, 2013). Challenges include discovering true dependencies, avoiding operational outages when rules are tightened, and maintaining visibility across heterogeneous environments.

Cloud-delivered security and SASE

Consolidating firewall, web gateway, and access functions in cloud services offers consistent policy and simpler management for distributed organizations. It also concentrates trust in the provider and creates new questions about data residency, service resilience, and the transparency of inspection. Research on the security consequences of concentrating enforcement in a few providers is limited.

Evolving encryption and privacy protocols

Encrypted DNS protocols such as DNS over HTTPS (Hoffman & McManus, 2018) prevent network devices from observing queries that were previously visible, and encrypted handshake extensions aim to conceal more connection metadata. The introduction of hybrid post-quantum key exchange may enlarge handshake messages and expose compatibility problems in middleboxes that make assumptions about traffic structure. Together these developments strengthen the case for endpoint- and identity-based telemetry and for cooperation between network and endpoint controls rather than reliance on passive network observation.

Artificial intelligence in security operations

Artificial intelligence may assist in rule analysis, log triage, and anomaly detection, and generative tools may help administrators write and review policies. The same caveats apply as in Section 5.7: performance must be validated in the target environment, adversaries may attempt to evade or poison models, and automated changes to enforcement policy require human oversight because an erroneous rule can either open a hole or cause an outage.

Automation and policy-as-code

Managing firewall policy through version control, peer review, automated testing, and continuous compliance checking directly targets the misconfiguration problem identified in earlier research (Wool, 2004, 2010). Expressing rules as code allows changes to be reviewed before deployment, rolled back if necessary, and tested against formal policy checks of the kind described in Section 6. This is one of the most promising practical

directions for reducing configuration-related failures.

A Framework for Evaluating Firewall Effectiveness

The review shows that a feature checklist is a poor proxy for protection. Effectiveness should instead be assessed along several complementary dimensions, summarized in Table 4. The framework combines configuration audits, coverage measurements, threat-informed testing, and operational indicators. It is proposed as a structure for organizational assessment and as a basis for future empirical validation, not as a validated instrument.

Table 4

Proposed seven-dimension framework for outcome-based firewall evaluation

Dimension	Example indicators	Evidence source	Suggested review cycle
1. Policy quality	Share of rules that are unused, redundant, shadowed, or overly permissive; rules without owner or justification; time since last review	Automated rule analysis; configuration audits	Quarterly
2. Coverage	Share of traffic and assets that pass through an enforcement point, including cloud and remote users	Asset inventory; flow logs; architecture review	Semi-annual
3. Visibility	Percentage of relevant encrypted traffic inspected or otherwise observable; log completeness	Traffic sampling; inspection statistics	Quarterly
4. Detection and prevention efficacy	Blocked versus missed techniques in adversary emulation mapped to ATT&CK	Red-team or breach-and-attack-simulation exercises	Annual or after major change

Dimension	Example indicators	Evidence source	Suggested review cycle
5. Resilience	Mean time to patch critical firewall vulnerabilities; exposure of management interfaces; recovery test results	Vulnerability scans; patch records; backup tests	Continuous / quarterly
6. Operational sustainability	Alert volume per analyst; false-positive rate; change-failure rate	SIEM and ticketing data	Monthly
7. Performance impact	Throughput and latency under the full inspection profile; number of inspection exemptions	Load testing; device telemetry	Annual or after capacity change

Note. Proposed by the authors on the basis of the literature reviewed in Sections 5 to 7.

Organizations can also position themselves on a simple four-level maturity scale (Table 5). The levels are descriptive: an organization may be at different levels for different dimensions, and progression is not automatically a matter of purchasing additional products.

Table 5

Illustrative maturity model for firewall management

Level	Characteristics
1. Ad hoc	Rules accumulate without review; vendor default settings; patching irregular; no measurement of effectiveness; management interfaces may be exposed.
2. Managed	Documented policy; named rule owners; periodic reviews; defined patch process; central log collection; basic segmentation.
3. Measured	Metrics tracked across the seven dimensions; governed selective TLS inspection; periodic adversary emulation; integration with identity and endpoint telemetry.
4. Optimized	Policy-as-code with automated testing; continuous verification of reachability; enforcement aligned with zero trust; feedback from incidents and exercises into policy.

Note. Illustrative model proposed by the authors; not empirically validated.

Combining periodic configuration audits with controlled adversary emulation offers a more realistic picture than vendor benchmarks alone. Where resources are limited, dimensions 1, 2, and 5 (policy quality, coverage, and resilience) typically deliver the largest improvement per unit of effort.

Recommendations

Recommendations for practitioners

Enforce default-deny and least-privilege rules; assign an owner and a business justification to each rule, set expiry dates for exceptions, and prune rule sets on a fixed schedule.

Treat firewalls and VPN gateways as high-value assets: restrict management access to dedicated networks, apply security patches

urgently, monitor device logs and integrity, and rehearse recovery from a compromised device.
Deploy internal segmentation so that a foothold does not equal broad access, starting with the most critical systems, backups, and management networks.
Adopt selective TLS inspection with clear governance, exclusions for sensitive categories, and secure certificate handling, and supplement it with DNS security and endpoint detection.
Use upstream DDoS mitigation rather than relying on enterprise edge devices to absorb large attacks.
Integrate firewall telemetry with identity, endpoint, and cloud logs in a central detection pipeline so that events can be correlated.
Move toward policy-as-code with peer review and automated verification to reduce configuration errors.
Validate defenses regularly through red-team or breach-and-attack-simulation exercises and use the results to update policy.

Priorities by organizational profile

The relative importance of these measures depends on organizational capacity. Table 6 offers a starting point.

Table 6

Suggested priority actions by organizational profile

Profile	Priority actions
Small organization or limited security staff	Default-deny policy; minimize exposed services; automatic or rapid patching of internet-facing devices; multi-factor authentication for remote access and administration; managed firewall or cloud-delivered service; off-line backups; log retention.
Mid-sized organization	All of the above, plus internal segmentation; quarterly rule review; centralized logging and alerting; selective TLS inspection; WAF for public applications; annual external testing.
Large or high-risk organization	All of the above, plus microsegmentation; policy-as-code and continuous verification; adversary-emulation programs; zero-trust network access; dedicated monitoring capability; formal metrics on the seven-dimension framework.

Note. Author synthesis; priorities should be adapted to sector, regulatory obligations, and risk assessment.

11.3 Recommendations for vendors, standards bodies, and researchers

Vendors should publish transparent descriptions of what inspection is performed under which configurations, provide secure-by-default settings, and disclose and remediate vulnerabilities quickly. Standards bodies could develop vendor-neutral test methodologies for encrypted-traffic environments. Researchers should prioritize replicable measurement, open data sets, and studies that connect firewall practices to incident outcomes.

Research Gaps and Future Work

Several areas merit further study:
Standardized, vendor-neutral benchmarks for NGFW efficacy under encrypted traffic and realistic traffic mixes.
Longitudinal measurement of misconfiguration in cloud-native rule sets, extending earlier work on traditional firewalls (Wool, 2004, 2010).
Robustness of machine-learning-based inspection against adversarial evasion and concept drift (Sommer & Paxson, 2010).
Quantitative comparison of segmentation and microsegmentation strategies for ransomware containment.
Security consequences of concentrating enforcement in a small number of cloud-delivered services, including resilience and transparency.
Empirical validation of multi-dimensional evaluation frameworks such as the one proposed here, including which dimensions best predict incident outcomes.
Behavioral and organizational studies of how change management, staffing, and ownership affect policy quality.
Effects of emerging protocol changes, including encrypted DNS, encrypted handshakes, and post-quantum key exchange, on in-path inspection.

13. Limitations of the Review

This review has several limitations. As a narrative review, it is subject to selection bias, and although the sources were chosen for relevance and credibility, another reviewer might have weighted them differently. Many of the most rigorous studies of firewall configuration date from earlier years, and equivalent evidence for cloud-native environments is less mature. The qualitative ratings in Table 2 and Figure 1 are an interpretive synthesis rather than experimental measurements. Threat statistics and vulnerability lists change rapidly, and the examples cited illustrate patterns rather than provide a complete inventory. The proposed evaluation framework and maturity model have not been empirically validated.

Conclusion

Firewalls remain a foundational control. They reliably block a large class of network-level threats and enable segmentation, egress control, and policy enforcement on which many other defenses depend. Their limits are equally clear. They cannot fully address encrypted, credential-based, insider, and supply-chain threats, they cannot absorb large volumetric attacks, and they are themselves attractive and sometimes vulnerable targets. The evidence reviewed here suggests that their value is determined mainly by disciplined configuration, rapid patching, adequate visibility, and integration into a broader zero trust and detection strategy, rather than by the sophistication of the product category alone.

For practitioners, the implication is to evaluate firewalls by measurable outcomes, to invest in policy hygiene and patching before adding features, and to treat the firewall as one layer within defense in depth. For researchers, the priority is to build the vendor-neutral, replicable evidence base that the field still lacks. Organizations that combine these practices are better placed to gain real protection from firewalls in an environment where the perimeter is porous and adversaries are adaptable.

References

Al-Shaer, E., & Hamed, H. (2004). Discovery of policy anomalies in distributed firewalls. In *Proceedings of IEEE INFOCOM 2004* (Vol. 4, pp. 2605-2616). IEEE.

Al-Shaer, E., Marrero, W., El-Atawy, A., & ElBadawi, K. (2009). Network configuration in a box: Towards end-to-end verification of network reachability and security. In *Proceedings of the 17th IEEE International Conference on Network Protocols (ICNP)* (pp. 123-132). IEEE.

Anderson, B., & McGrew, D. (2016). Identifying encrypted malware traffic with contextual flow data. In *Proceedings of the 2016 ACM Workshop on Artificial Intelligence and Security (AISec)* (pp. 35-46). ACM.

Anderson, B., Paul, S., & McGrew, D. (2018). Deciphering malware's use of TLS (without decryption). *Journal of Computer Virology and Hacking Techniques, 14*(3), 195-211.

Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., ... Zhou, Y. (2017). Understanding the Mirai botnet. In *Proceedings of the 26th USENIX Security Symposium* (pp. 1093-1110). USENIX Association.

Bellovin, S. M. (1989). Security problems in the TCP/IP protocol suite. *ACM SIGCOMM Computer Communication Review, 19*(2), 32-48.

Bellovin, S. M. (1999). Distributed firewalls. *;login:, 24*(Special Issue), 37-39.

Cheswick, W. R., Bellovin, S. M., & Rubin, A. D. (2003). *Firewalls and Internet security: Repelling the wily hacker* (2nd ed.). Addison-Wesley.

Cisco Talos. (2024). *ArcaneDoor: New espionage-focused campaign found*

targeting perimeter network devices. Cisco Talos Intelligence Group.

Cybersecurity and Infrastructure Security Agency. (n.d.). *Known exploited vulnerabilities catalog*. https://www.cisa.gov/known-exploited-vulnerabilities-catalog

Cybersecurity and Infrastructure Security Agency. (2020). *Alert AA20-352A: Advanced persistent threat compromise of government agencies, critical infrastructure, and private sector organizations*. CISA.

Durumeric, Z., Ma, Z., Springall, D., Barnes, R., Sullivan, N., Bursztein, E., Bailey, M., Halderman, J. A., & Paxson, V. (2017). The security impact of HTTPS interception. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*. Internet Society.

European Union Agency for Cybersecurity. (2023). *ENISA threat landscape 2023*. ENISA.

Gilman, E., & Barth, D. (2017). *Zero trust networks: Building secure systems in untrusted networks*. O'Reilly Media.

Gouda, M. G., & Liu, A. X. (2007). Structured firewall design. *Computer Networks, 51*(4), 1106-1120.

Hamed, H., & Al-Shaer, E. (2006). Taxonomy of conflicts in network security policies. *IEEE Communications Magazine, 44*(3), 134-141.

Handley, M., Paxson, V., & Kreibich, C. (2001). Network intrusion detection: Evasion, traffic normalization, and end-to-end protocol semantics. In *Proceedings of the 10th USENIX Security Symposium*. USENIX Association.

Hoffman, P., & McManus, P. (2018). *DNS queries over HTTPS (DoH)* (RFC 8484). Internet Engineering Task Force. https://doi.org/10.17487/RFC8484

Hu, H., Ahn, G.-J., & Kulkarni, K. (2012). Detecting and resolving firewall policy anomalies. *IEEE Transactions on Dependable and Secure Computing, 9*(3), 318-331.

Hutchins, E. M., Cloppert, M. J., & Amin, R. M. (2011). Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare & Security Research, 1*, 80-106.

Ingham, K., & Forrest, S. (2002). *A history and survey of network firewalls* (Technical Report TR-CS-2002-37). University of New Mexico.

Kim, H., & Feamster, N. (2013). Improving network management with software defined networking. *IEEE Communications Magazine, 51*(2), 114-119.

Kindervag, J. (2010). *No more chewy centers: Introducing the zero trust model of information security*. Forrester Research.

Kolias, C., Kambourakis, G., Stavrou, A., & Voas, J. (2017). DDoS in the IoT: Mirai and other botnets. *Computer, 50*(7), 80-84.

Liu, A. X., & Gouda, M. G. (2008). Diverse firewall design. *IEEE Transactions on Parallel and Distributed Systems, 19*(9), 1237-1251.

McHugh, J. (2000). Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by Lincoln Laboratory. *ACM Transactions on Information and System Security, 3*(4), 262-294.

Mirkovic, J., & Reiher, P. (2004). A taxonomy of DDoS attack and DDoS defense mechanisms. *ACM SIGCOMM Computer Communication Review, 34*(2), 39-53.

MITRE. (n.d.). *MITRE ATT&CK*. https://attack.mitre.org

National Institute of Standards and Technology. (2024). *The NIST Cybersecurity Framework (CSF) 2.0* (NIST CSWP 29). https://doi.org/10.6028/NIST.CSWP.29

National Vulnerability Database. (n.d.). *CVE-2021-44228* [Vulnerability record]. NIST. https://nvd.nist.gov/vuln/detail/CVE-2021-44228

National Vulnerability Database. (n.d.). *CVE-2022-42475* [Vulnerability record]. NIST. https://nvd.nist.gov/vuln/detail/CVE-2022-42475

National Vulnerability Database. (n.d.). *CVE-2024-3400* [Vulnerability record]. NIST. https://nvd.nist.gov/vuln/detail/CVE-2024-3400

OWASP Foundation. (2021). *OWASP Top 10: 2021*. https://owasp.org/Top10/

Paxson, V. (1999). Bro: A system for detecting network intruders in real-time. *Computer Networks, 31*(23-24), 2435-2463.

Ptacek, T. H., & Newsham, T. N. (1998). *Insertion, evasion, and denial of service: Eluding network intrusion detection*. Secure Networks, Inc.

Rescorla, E. (2018). *The Transport Layer Security (TLS) protocol version 1.3* (RFC 8446). Internet Engineering Task Force. https://doi.org/10.17487/RFC8446

Roesch, M. (1999). Snort: Lightweight intrusion detection for networks. In *Proceedings of the 13th USENIX Conference on System Administration (LISA '99)* (pp. 229-238). USENIX Association.

Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero trust architecture* (NIST Special Publication 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207

Scarfone, K., & Hoffman, P. (2009). *Guidelines on firewalls and firewall policy* (NIST Special Publication 800-41 Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-41r1

Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP)* (pp. 108-116). SciTePress.

Sherry, J., Hasan, S., Scott, C., Krishnamurthy, A., Ratnasamy, S., & Sekar, V. (2012). Making middleboxes someone else's problem: Network processing as a cloud service. *ACM SIGCOMM Computer Communication Review, 42*(4), 13-24.

Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. In *Proceedings of the 2010 IEEE Symposium on Security and Privacy* (pp. 305-316). IEEE.

Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. In *Proceedings of the 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications* (pp. 1-6). IEEE.

Verizon. (2024). *2024 data breach investigations report*. Verizon Business.

Ward, R., & Beyer, B. (2014). BeyondCorp: A new approach to enterprise security. *;login:, 39*(6), 6-11.

Wool, A. (2004). A quantitative study of firewall configuration errors. *Computer, 37*(6), 62-67.

Wool, A. (2010). Trends in firewall configuration errors: Measuring the holes in Swiss cheese. *IEEE Internet Computing*, *14*(4), 58-65.

Zargar, S. T., Joshi, J., & Tipper, D. (2013). A survey of defense mechanisms against distributed denial of service (DDoS) flooding attacks. *IEEE Communications Surveys & Tutorials*, *15*(4), 2046-2069.

IEER
Institute for Excellence in Education & Research