Connection context verification as a basis for network access control in zero-trust architectures*

Roman Syrotynskyi[1,*,†], Ivan Tyshyk[1,†], and Bohdan Zhurakovskiy[2,†]

[1] Lviv Polytechnic National University, Information Security Department, 12 Stepan Bandera str., 79000 Lviv, Ukraine
[2] National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute," 37 Peremogy ave., 03056 Kyiv, Ukraine

Abstract

The need to adopt new approaches and practices for organizing and governing access in network infrastructures is driven by the growing disparity between the demands of modern security standards and the limitations of network security controls that primarily operate at OSI Layers 3–4. This paper examines security models promoted by contemporary standards and industry best practices for protecting information infrastructures, and investigates practical pathways for implementing them using widely available market solutions and network access control measures. The proposed methodology integrates the Zero Trust principles described in NIST SP 800-207 with capabilities commonly offered in next-generation firewall vendor portfolios, complemented by the author's methods for integrating heterogeneous security systems to enrich firewall policy decisions with network-access context. The approach supports alignment with Zero Trust principles while preserving operational quality and high network performance, without surpassing acceptable thresholds for total cost of ownership and operation. Key components and recurring design patterns required for such a security infrastructure are identified.

The scientific novelty of the work is a shift in the network access control paradigm—from policies built around a corporate node's address to policies centered on user-driven authorization combined with verification of the requesting device's security posture. The paper proposes a set of contextual attributes for security policies and presents effective methods for structuring host access levels within firewall configurations. By collecting and augmenting policy enforcement points with network context, the approach provides the flexibility and technical foundation necessary to uphold Zero Trust principles in corporate security architectures. At the same time, it introduces trade-offs, including increased operational complexity and cost, as well as dependencies on external systems that can influence network performance and potentially enlarge the overall compromise surface of the security stack.

Keywords

Connection context, NGFW, ZeroTrust, user identity mapping, device/host security posture, dynamic address group.

1. Introduction

In the current environment of rapidly evolving cyber threats, the traditional perimeter-based security model is losing its effectiveness. The widespread use of cloud services, mobile workplaces, and interconnected corporate networks means that any implicit trust within the network becomes a significant risk. To address these challenges, organizations are adopting Zero Trust Architecture (ZTA)—a security model that abandons the concept of "trusted zones" and is based on the principle of "never trust, always verify."

Network access control mechanisms and firewall policies play a key role in implementing ZTA principles. Whereas firewalls were previously based primarily on the analysis of IP addresses, ports, and application-layer protocols, they have now become critical decision points for granting access in microsegmented environments. Within the Zero Trust model, firewall policies cannot remain static; instead, they must dynamically adapt to user identity, device posture, workload sensitivity, and other contextual factors.

* CPITS 2026: Cybersecurity Providing in Information and Telecommunication Systems, February 28, 2026, Kyiv, Ukraine
* Corresponding author.
† These authors contributed equally.
roman.m.syrotynskyi@lpnu.ua (R. Syrotynskyi); ivan.y.tyshyk@lpnu.ua (I. Tyshyk); zhurakovskybiyu@tk.kpi.ua (B. Zhurakovskiy)
0009-0002-6280-3290 (R. Syrotynskyi); 0000-0003-1465-5342 (I. Tyshyk); 0000-0003-3990-5205 (B. Zhurakovskiy)

In this paradigm, contextual policies implemented in next-generation firewalls (NGFWs) gain particular importance. These policies go beyond static rule sets by incorporating attributes such as user roles, device security posture, geographic location, application type, and even behavioral indicators into access control decisions. By leveraging contextual data, modern NGFWs provide greater granularity, flexibility, and risk reduction compared to overly permissive or outdated rules. Context-aware policies thus become a necessary condition for maintaining a balance between security and usability in large-scale, distributed, and hybrid infrastructures.

However, expanding firewall functionality to support monitoring and validation of network access context introduces new challenges and tasks, including additional system dependencies, increased technological complexity, substantial operational overhead, and emerging cybersecurity risks.

Problem Statement. Traditional network access control systems have historically relied on static firewall rules defined by IP addresses and TCP/UDP port-based service protocols. While this approach enables basic traffic segmentation, it is no longer sufficient in modern, large-scale, and highly dynamic corporate networks. Within the context of Zero Trust Architecture, static rules introduce several limitations: they are overly generalized, lack sufficient flexibility, and often result in excessive access privileges. This creates the risk of uncontrolled privilege expansion to critical resources and complicates security maintenance in environments where users, devices, and services continuously change their interaction parameters.

Purpose of the Study. The objective of this article is to examine the role of contextual firewall policies and the specifics of their application within Zero Trust Architecture. In particular, the study explores types of additional security policy attributes used to enhance access control, technologies that enable the collection and integration of contextual data into NGFW operations, and strategies and approaches for managing and optimizing context-aware rules in scalable network environments.

2. Analysis of Recent Research and Publications

Adherence to the principles of Zero Trust Architecture necessitates the use of more sophisticated methods and tools for network traffic management. Researchers emphasize the limitations of static access policies in modern cloud-based and distributed environments, noting that such mechanisms cannot adapt to changes in user behavior or device state. This creates a need for more dynamic and fine-grained access control mechanisms [1]. Context-aware access control enables the implementation of adaptive policies based on real-time risk assessment, often leveraging telemetry from identity and security platforms.

Recent architectural approaches demonstrate how contextual data can be integrated into ZTA. For example, Routray and Bera (2025) proposed a Zero Trust Adaptive Authorization Architecture (ZTAAC) that employs Ciphertext-Policy Attribute-Based Encryption (CP-ABE), allowing access only when the correct combination of static and dynamic attributes—such as device posture, user location, and behavioral patterns—is satisfied [2].

Similarly, Ahmadi (2025) introduced an autonomous threat segmentation model that applies machine learning to analyze login patterns, device type, geolocation, and access time in order to assign dynamic risk levels. Based on these assessments, the system automatically adjusts access permissions, minimizing insider threats and lateral movement within the network [3].

Practical implementations of context-aware ZTA approaches have also been demonstrated in enterprise environments. For instance, within cloud-based .NET environments, conditional access policies have been implemented using Azure Active Directory and OAuth 2.0, taking into account user identity, device compliance with security policies, and geographic location. These policies adapt to threats in real time through integration with security platform telemetry [4].

In federated systems, contextual information is collected from multiple data sources. In the Zero Trust Federation (ZTF) model proposed by Hirai et al., a centralized Context Aggregation Point

(CAP) gathers data from sources such as RADIUS and Mobile Device Management (MDM) systems. This enables more accurate access decisions based on aggregated contextual information [5].

The integration of artificial intelligence into Zero Trust Architecture enables more flexible access governance. For example, Oluoha and Odeshina proposed an AI-based framework for Continuous Access Governance that analyzes behavioral patterns, device trust levels, and session risks, automating role management and access verification in compliance with GDPR and NIST requirements [6].

To systematize knowledge in the field of contextual security, Xiao et al. conducted a comprehensive study highlighting the intersections between context-aware, risk-aware, and Zero Trust policies. The authors argue that the use of dynamic attributes—such as user location, behavior, and device state—significantly improves the accuracy of access control decisions [7].

At the same time, several challenges remain. Integrating contextual verification into legacy systems requires significant architectural changes and data source unification [8]. Privacy concerns related to the collection of behavioral and environmental data also arise, potentially leading to false positives or degradation of user experience [9].

3. Materials and Methods

3.1 Network connection context overview. Traditionally, network access is described using fields such as IP addresses, protocol, port, and actions allow or deny. This approach to regulating network access operates primarily at Layers 3 and 4 of the OSI model and, in practice, grants access not to a specific host or system, but rather to whoever sends packets with matching values in the "source IP," "destination IP," and "protocol/port" fields. Early firewall technologies were based on these principles, which was largely acceptable at the time. However, organizing network access control in accordance with Zero Trust principles using only basic conditions (IP and port fields) is insufficient to ensure the security of modern network architectures due to limited flexibility, excessive static behavior, lack of sender authentication, and restricted visibility.

Understanding the context of a connection significantly expands access management capabilities. By incorporating additional contextual attributes, security policies can more precisely define whether specific access should be allowed or denied. Enriching firewall policies with attributes that describe the context of a network connection opens new possibilities for access control, primarily by enabling more flexible connectivity models. Additional validation checks allow access to be restricted in a more targeted manner, reducing the risk of IP spoofing and other attacks that rely on mimicking permitted access.

Integrating identity verification into firewall policies eliminates the need to rely solely on the network address of the host to which access is granted or denied. This fundamentally changes the concept of network traffic access control. When the "identity" attribute is applied, firewall policies control access for a specific user or service rather than for the network segment in which they reside. This shift in approach provides several advantages:

- the ability to distinguish the connection of a specific user or system from other members of the same network;
- continuous identity verification ensures compliance with Zero Trust principles;
- simplified access management in dynamic environments where users or services may appear in different network segments;
- a conceptual change in dynamic access control, from adding a host to an allowed network to adding a user to an authorized identity provider group (or removing them).

Application-level inspection within firewall rules, either in addition to or as a replacement for traditional port/protocol matching, enables more precise control over traffic that uses the same port but corresponds to different application types. For example, application identification allows different types of traffic over TCP port 443 to be distinguished and controlled independently.

Contextual attributes used in firewall rules can be grouped into the following categories:

User and Identity Context. This category includes attributes that characterize the access subject, such as user role, organizational unit membership, and authentication method. Integration of these attributes is typically achieved through directory services (Active Directory, LDAP) or Single Sign-On (SSO) systems using protocols such as SAML or OAuth. This enables policy enforcement based on functional user roles and access control that reflects the user's operational context.

Device and Endpoint Context. This context covers characteristics of the devices from which access is initiated, including device type, operating system version, and compliance with security policies (e.g., patch level or presence of antivirus protection). Integration is commonly achieved through Mobile Device Management (MDM) systems and Endpoint Detection and Response (EDR) solutions. Considering these parameters allows access from vulnerable or misconfigured devices to be blocked.

Network and Session Context. This type of context relates to network connection parameters, such as source and destination zones, subnet, VPN connection status, and client geolocation. These attributes can be used, for example, to restrict access to trusted subnets only or to users authenticated through a corporate VPN.

Application and Data Context. This category includes attributes describing application type, service, or data sensitivity level. It enables access policies based on information resource classification, such as allowing interaction with confidential documents only through approved corporate services (e.g., SharePoint). This approach integrates network security with data protection mechanisms.

Environmental Context. Environmental context accounts for dynamic external factors, including time of day, day of the week, physical user location, and detected anomalies in device behavior. Applying such attributes enables restrictions on critical operations, for example, blocking administrative access to cloud resources outside business hours or from atypical geolocations.

3.2 NGFW Technologies for Supporting Contextual Policies. Technologies for collecting and validating contextual information extend beyond what can be conveyed using standard IP packet fields. Without additional mechanisms, a firewall cannot determine the context of a network connection. The technical means required to implement contextual validation depend on the type of context involved and vary accordingly.

Building security policies based on identity verification requires solving several nontrivial challenges, including identifying the user initiating the connection and verifying the user account's membership in identity directory groups. Since firewalls operate at the network level, a user identification mechanism is required to associate network connections with user identities. Modern next-generation firewalls provide integration technologies with identity providers and other systems that enable mapping between IP addresses and user accounts.

As a result, when a firewall policy specifies a particular user account, the firewall can determine from which IP address that user will connect and dynamically allow access only for that IP address. The core task lies in building a user–ID mapping table and delivering it to the firewall for use in security policies. Mechanisms are also required to retrieve group information from directory services and verify user account membership in those groups.

User identity–based policies operate by associating a user identifier with the user's current IP address or session context. User identifiers are typically obtained from centralized authentication systems or directories such as Active Directory, LDAP, SAML-based Single Sign-On, or cloud identity providers such as Azure Active Directory or Okta (see Figure 1). This mapping enables the firewall to enforce access rules at the level of individual users or groups rather than network segments.

The core technology underlying this implementation is the user identity mapping mechanism, which can be established in several ways.

- Agent-based log monitoring: a dedicated user identification agent monitors authentication logs (e.g., Kerberos or NTLM events) on domain controllers to associate IP addresses with user logins.

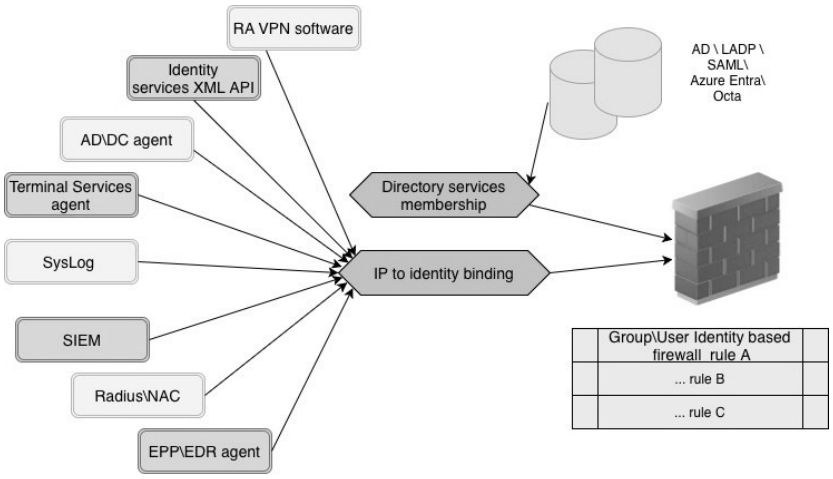

Figure 1. Identity-based access control of network traffic

- Syslog/API-based identity data collection: firewalls receive syslog data from VPN concentrators, RADIUS servers, or Network Access Control (NAC) systems that contain authentication events.

- Endpoint-based identity reporting: firewall agents or endpoint protection platforms (EPP/EDR) directly report user identity and posture information to the firewall.

- Authentication portals and SAML assertions: in guest environments or on unmanaged devices, firewalls may redirect users to authentication portals or use SAML assertions to identify user credentials.

Verification of endpoint posture and security status is a critical factor in making decisions about granting or denying network access. Not only statically defined rules and conditions are important, but also the firewall's ability to dynamically react to changes in endpoint posture. To enable the use of device context in firewall policies, the following technologies may be applied:

- Endpoint Detection and Response (EDR) / Extended Detection and Response (XDR). In these solutions, EDR/XDR agents continuously monitor endpoint behavior, file integrity, process activity, and potential compromise indicators. Leading vendors such as CrowdStrike, Microsoft Defender for Endpoint, and Palo Alto Cortex XDR provide APIs or generate alerts that can be consumed by firewalls or SOAR platforms.

- Mobile Device Management (MDM) / Unified Endpoint Management (UEM). MDM solutions such as Jamf, Intune, or VMware Workspace ONE maintain compliance and configuration data that can be used for access control decisions. Firewall integration is often achieved via tags (e.g., "compliant," "non-corporate"), device identifiers, or API-based status queries. This enables policies such as: "Block devices not enrolled in MDM from accessing sensitive applications."

- Native NGFW endpoint integration. Some NGFW vendors collect endpoint data directly via proprietary agents. For example, Palo Alto GlobalProtect HIP gathers information about operating system version, antivirus status, disk encryption, and more; Fortinet FortiClient reports device compliance and posture information to FortiGate; Cisco AnyConnect combined with Cisco ISE provides firewall posture data via ISE and pxGrid.

- SOAR / Context aggregators. Security orchestration tools (e.g., Cortex XSOAR, Splunk SOAR, Tines) can normalize data from multiple sources and push contextual tags to firewalls via APIs or syslog. These tags (e.g., "high_risk_device") can then be consumed by dynamic address groups (Palo Alto) or Zero Trust tags (Fortinet).

Scenario: a device with outdated antivirus software is flagged by CrowdStrike. The sequence of actions is as follows:

- CrowdStrike EDR sends an alert to XSOAR.
- XSOAR maps the IP address with the tag AV_outdated.
- The tag is transmitted to the Palo Alto NGFW via API.

- NGFW policy: DENY access to internal applications IF the source IP address belongs to the DAG: AV_outdated.

- Access to the device remains blocked until the antivirus software is updated and the tag is removed.

Application-aware filtering is a core capability of next-generation firewalls (NGFWs), enabling the enforcement of security policies based on application-layer (Layer 7) information rather than solely on IP addresses, ports, or protocols. This technology is based on Deep Packet Inspection (DPI), which analyzes packet payloads to accurately identify applications regardless of the port used or the presence of encryption. In addition, protocol decoders and heuristic traffic analysis are employed to recognize application behavior and detect evasion techniques such as tunneling or port hopping. Some firewalls perform SSL/TLS decryption to inspect encrypted traffic, providing visibility into applications operating over HTTPS. Regularly updated application signatures supplied by vendors are matched against traffic flows to classify thousands of known applications. Vendors such as Palo Alto Networks (App-ID), Fortinet (App Control), and Cisco (NGFW Snort/NGIPS) implement proprietary mechanisms to support this functionality. Application-aware filtering also enables the enforcement of identity- and role-based policies by restricting specific actions within an application (e.g., allowing Facebook chat but blocking file uploads). Integration with cloud application databases and real-time threat intelligence further improves the accuracy of web and SaaS traffic identification. Ultimately, application-aware filtering enhances policy precision and reduces the attack surface by allowing organizations to define access based on application usage rather than mere connectivity [10].

Time and location context. Policies may restrict access based on time of day, geolocation, or network zone. Examples include blocking administrative access outside business hours or denying VPN logins from untrusted countries. Time- and location-based contextual policies enable the implementation of dynamic ZTA rules beyond static access definitions [11].

Beyond traditional firewall generations and commonly used contextual inputs—such as user identity, device posture, and application awareness—new technologies have emerged to support advanced contextual policy enforcement. One such development is the use of model-driven access control frameworks, such as GemRBAC+CTX, which allow spatial–temporal attributes (e.g., time, location, environmental state) to be explicitly defined within access control policies. These systems formalize contextual constraints using models and logical languages such as OCL, enabling enforcement mechanisms to reason about complex scenarios without relying exclusively on static parameters [12].

Another emerging approach involves blockchain-enhanced access control, where dynamic attributes such as time or location are cryptographically enforced through smart contracts, enabling distributed trust and protection against unauthorized access. This makes policy enforcement more transparent and resilient in cloud environments [13]. Collectively, these advanced technologies highlight a shift toward contextual decentralization, allowing policies to evolve in real time and leverage richer, multidimensional inputs than previously possible.

3.2 Problems and Limitations of Applying Contextual Policies. Undoubtedly, the use of contextual policies increases the flexibility of network access management and enhances the security of information infrastructures. The transition to the ideology of “access for the user rather than for the network in which the user resides” fundamentally changes the firewalling paradigm in corporate networks.

The need to synchronize access rules depending on network entry points and the task of differentiating one user’s connection from others at the firewall rule level—especially in networks with dynamic addressing—becomes largely irrelevant. Instead, a single set of network rules can be defined for a specific user account or group of accounts.

However, this advanced functionality is complex and costly to implement. Only market-leading vendors provide comprehensive support for contextual policies out of the box and offer built-in integration mechanisms and tools for collecting contextual information within their product

portfolios. Moreover, network service operation increasingly depends on external systems, which expands the potential landscape of failures and compromise points within information infrastructures.

The complexity of configuring, maintaining, and analyzing firewall security policy behavior also increases. Enriching policies with additional contextual attributes leads to the problem of overly complex rules. Such rules may include a large number of conditions, attributes, or exceptions, making them difficult to understand and maintain. While this complexity may be necessary for precise access control, excessively complex rules increase the risk of misconfiguration, security gaps, and performance degradation.

For example, consider a corporate firewall rule designed to allow access to a confidential internal web server (SecretServer) only for employees using managed devices, located in specific office locations, during business hours, holding certain roles (e.g., R&D or IT), and connecting from specific IP address ranges.

Below is an example of an overly complex rule:

- If the user is authenticated with Corp Directory and belongs to the "OPS" or "IT" group
- If the user is accessing from one of the following offices (Rivne, Cykiv, Chernivtsi)
- If the client device has a valid endpoint security posture (AV installed, OS version xyz or higher)
- If the source IP is one of the following (10.20.10.25, 10.20.20.25, 10.20.30.45, etc.)
- If the access request is between 8 AM to 8 PM local time initiated from office network
- If the device is connected via the corporate VPN
- If the user hasn't failed more than 5 authentication attempts in the past 3 hours

The implementation of excessively complex contextual security policies leads to a number of problems and operational challenges, including the following:

- Excessive number of conditions. Multiple attributes are evaluated simultaneously (e.g., user role, device security posture, time, source IP address, VPN configuration, and failed login attempts). This level of complexity increases the risk of misconfiguration and makes it difficult to verify whether all conditions are correctly defined and consistently enforced.

Increased maintenance overhead. As the lists of IP address ranges, office locations, or user roles grow, policy management becomes increasingly difficult, particularly in environments with frequent network changes or evolving user roles.

Troubleshooting complexity. When access for an authorized user is blocked and traffic does not match a large, cumbersome security policy, identifying the root cause can be challenging. It may be unclear whether the denial is due to device posture, source IP address, failed authentication attempts, or office location.

Another example of a problematic scenario is a configuration containing multiple nested rules for specific use cases. For instance, consider a rule in which an administrator attempts to define different access levels for the same host or user based on a combination of roles, device type, and time of day:

***Rule 1**: Allow "OPS" users to access Server if:*
- Device is a Windows laptop or desktop
- Access is from Rivne or Cykiv offices
- Between 8 AM to 8 PM local time
- If they are part of "OPS" Active Directory group

***Rule 2**: Allow "DevOps" users to access SecretServer if:*
- Device is a Windows laptop or desktop
- Access is from any office in Ukraine office

- Between 6 AM to 8 AM or 7 PM to 9 PM
- If they are part of "DevOps" Active Directory group
- If account have MFA enabled

***Rule 3**: Allow "Manager" group to access SecretServer if:*
- Device is a Macbook
- Access is from Lviv office
- During any time of day
- If they are part of "Management" Active Directory group

This approach to rule configuration introduces challenges similar to the effects of excessive configuration accumulation, including the following:

Multiple nested rules. Each rule introduces numerous conditions and specific access time windows, resulting in a complex and overlapping rule set. The firewall may be required to evaluate multiple conditions before granting access, which can negatively affect performance.

Increased risk of rule overlap. Such nested rules can lead to redundancy and potential overlaps, where different roles fall under conflicting conditions.

Ambiguity in rule enforcement. In the presence of confirmation conflicts between rules (e.g., differing time constraints or device types), it can be difficult to determine which rule takes precedence and how they interact. The practice of creating exceptions for individual users on top of generally accepted rules significantly increases configuration complexity, places additional processing load on the firewall—potentially resulting in performance degradation—scales poorly, and becomes difficult to analyze during troubleshooting and incident response.

3.3 Designing Effective Contextual Policies. To derive real value from the adoption of advanced networking and security technologies, it is essential to strike a balance between extensive configurations and complex cyber systems on the one hand, and concise rule sets with a limited number of dependencies on external systems on the other. An effective security system is one that does not degrade the overall performance of the network infrastructure and can be maintained by the existing technical staff within acceptable operational timeframes.

In the case of contextual policies, a key challenge lies in identifying the most effective combinations of contextual attributes and the rules built upon them. A well-balanced policy is one that enforces Zero Trust principles while remaining manageable, clear, and efficient. It should be sufficiently precise to ensure policy compliance without becoming overly complex, and it should focus on critical contextual elements such as user identity, device security posture, and time-based access, avoiding unnecessary exceptions.

The fundamental principles for designing contextual policies are clarity, flexibility, and structural consistency. A policy must scale effectively: as the organization grows, it should be easily adaptable to new user groups or resources by modifying the associated user or resource group (for example, replacing “Human Resources” with “Finance” or “Administrator”). Policies should not degrade firewall performance or negatively affect network quality characteristics. Processing traffic with contextual attribute validation may increase latency due to the involvement of additional systems. Therefore, the number of rules associated with a single user should be kept to a minimum, while still adhering to the principle of least privilege.

Maintainability: as organizational requirements evolve, the existing set of contextual rules should allow changes to conditions (e.g., adjusting time windows or introducing new device requirements) without overcomplicating the policy structure itself. Complex hierarchical layers of contextual policies within firewalls should be avoided. Likewise, the practice of defining numerous exceptions to generally applicable policies is discouraged, as such non-standardized rules are difficult to maintain and scale.

The principle of least privilege should serve not only as a driver for introducing additional contextual attributes into a policy, but also as an indicator of when a policy is already sufficiently

precise and does not require further refinement. Overly saturated policies create significant troubleshooting challenges, particularly in situations where traffic fails to match a rule and the specific unmet condition must be identified.

For each contextual attribute evaluated by a policy, a corresponding validation mechanism must exist. This is essential for diagnostics when a policy does not match traffic, as firewall logging typically occurs only when traffic successfully matches a rule [14]. If access conditions defined by contextual attributes are not met, the policy is not triggered and no log entry is generated, complicating root-cause analysis.

By using contextual firewall policies, it is possible to design a structured hierarchy of access levels for users or systems based on the degree of compliance with policy conditions. This approach enables a gradual reduction of access privileges as contextual compliance decreases. However, this often leads to the creation of so-called nested rules, where policies overlap and differ only slightly in their conditions and outcomes. The use of multi-level nested rules is not recommended, as it significantly complicates ongoing operational support and may adversely affect network performance.

A sufficiently effective access control system can be achieved using a three-level access model. This model defines three levels of access, with optional user notification, depending on the degree to which policy conditions are met:

- Full access
- Restricted access
- No access

The use of templates and standardized rule structures with variables for roles, devices, applications, and network segments significantly reduces duplication, simplifies auditing, and improves maintainability. In addition, policies should be adaptive and capable of responding in real time to contextual changes, such as device non-compliance, anomalous user behavior, or elevated threat levels identified by integrated threat intelligence or behavioral analytics systems.

Maintaining a balance between security and usability is critical. Overly restrictive policies can hinder legitimate operations and encourage workarounds, while insufficiently detailed rules may permit unauthorized access. Modular and concise policy structures help minimize complexity, reduce the likelihood of misconfiguration, and improve firewall performance. Integration with external context sources—such as IAM, MDM/EDR, threat intelligence platforms, and SIEM/UEBA systems - enables dynamic and well-informed policy enforcement.

Table1.
An example of a balanced contextual policy

Aspect	Condition	Action
User role	Finance	Permit
Endpoint status	Corporate, updated, software installed	
Source data	VPN / Trusted network	
Schedule	07:00 – 11:00	
Destination data	CMDB	
Application	SSL, SSH	
	All counted cinditions	
	Any other case	Deny

This table demonstrates a clear, manageable, and controlled policy that incorporates multidimensional context and can dynamically adapt to changes in device posture or access attempts outside business hours, thereby reflecting Zero Trust principles. The application of role-based access control using user groups, combined with device posture validation, ensures compliance with Zero Trust requirements.

3.4 Technological Solutions Using Contextual Policies. The application of contextual and dynamic policy technologies enables the coverage of a class of technical and business challenges that extend beyond the capabilities of classical firewalls and traditional approaches to network access management based on static IP policies. This section outlines several technical tasks that arise in corporate information infrastructures under contemporary security challenges.

Automated management of employee network access. Granting or revoking access to corporate resources is an integral part of employee onboarding and offboarding processes. In large organizations, this is a highly dynamic and continuous activity. User identity verification technologies implemented in next-generation firewalls enable access control to be managed using Active Directory objects within security policies instead of IP addresses. This approach provides the advantage of fine-grained access control at the level of individual users or user groups. By leveraging Active Directory groups in security policies, the point at which access is granted or revoked is moved outside the firewall itself, simplifying access lifecycle management and reducing administrative overhead.

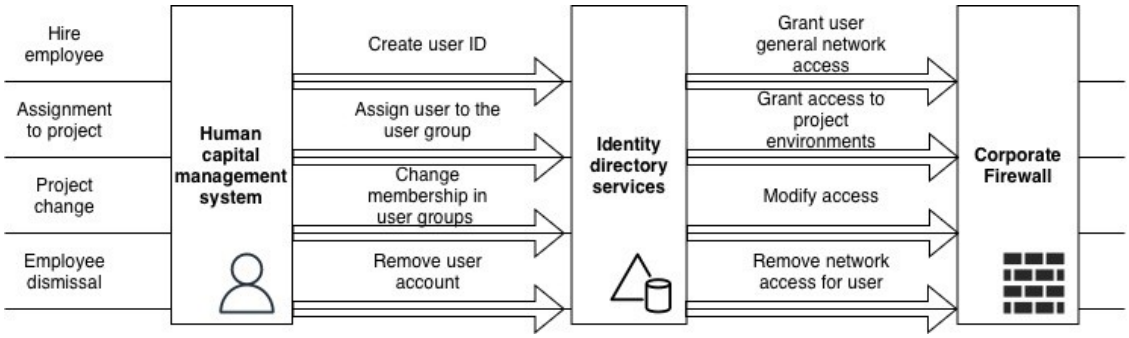

Figure 2. Organization of network access management in the user life cycle

Thus, when a user object is created in Active Directory and placed into a specific group that is already mapped within the firewall and referenced in a security rule, access for the new user is granted automatically without any modification to the rule configuration. This approach shifts the decision point for granting or denying network access outside the network firewall itself [15]. As a result, teams responsible for specific business processes can avoid delays and manage user access independently, without reliance on infrastructure operations teams.

An example of automated network access management using integrations between a Human Capital Management (HCM) system, a corporate directory service, and a firewall is illustrated in Figure 2.

Modern information infrastructures may rely on various dynamic environments located outside internal networks, for example, on the Internet. Such services may include cloud service providers as well as security vendors that specialize in providing lists of IP addresses or domain names with low reputational trust. These IP address lists are highly dynamic and may change daily, making manual access configuration labor-intensive and inefficient.

Dynamic access list technologies enable firewalls to retrieve lists of IP addresses, networks, or DNS names from external sources and automatically track changes at defined intervals. Examples of external and internal environments with dynamic addressing are shown in Figure 3. By configuring the firewall to consume these external feeds, access policies for dynamic environments remain up to date despite ongoing changes.

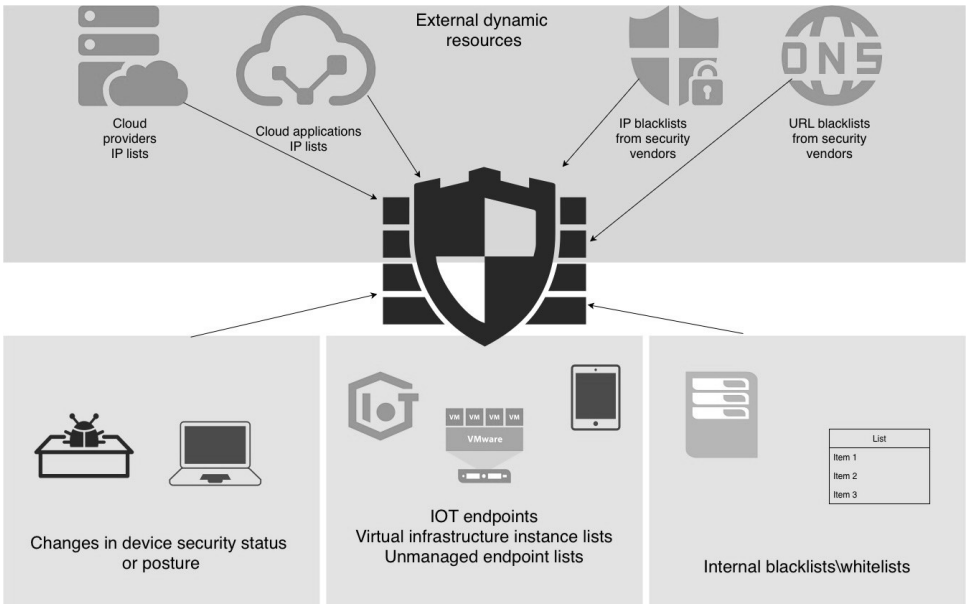

Figure 3. Corporate infrastructure dynamic environments

When dynamic access lists are applied, a time-driven control logic is used. Changes to the list itself occur only at the moment the list is retrieved from an external source, which happens periodically according to a configured time interval and is not event-driven, except in the case of manual triggering.

An enterprise cybersecurity infrastructure is capable of monitoring data flows, detecting anomalies or deviations, and responding to them promptly. Access to sensitive data or protected systems is granted only when clearly defined conditions are met [16, 17]. If, for any reason, network nodes no longer satisfy these conditions, a mechanism is required for the immediate correction of their access to sensitive resources. This situation creates the need for automatic and dynamic access correction for nodes whose cybersecurity posture has changed. Examples of such changes include the emergence of vulnerabilities on endpoints, removal of devices from corporate management tools, or failure by a user to complete mandatory security training. Consequently, there arises a requirement for event-driven access correction logic.

Examples of technologies that enable the dynamic population of firewall address groups with IP addresses and the assignment of specific tags include mechanisms such as Dynamic Address Groups (DAGs). These allow the creation of IP address groups associated with unique tags that can be dynamically populated by external systems using REST APIs as a result of specific events.

Consider a scenario illustrating the use of DAGs: when an antivirus system detects malware on an endpoint, it registers an event that adds the host's IP address to a dynamic group tagged "virus_infected". On the firewall, a security policy is configured to include the dynamic address group "virus_infected" and is automatically applied to all IP addresses registered within it. As a result, the system can automatically adjust the access of an endpoint or user when its cybersecurity posture changes, without human intervention.

This technology enables continuous monitoring of changes in a network node's context and supports immediate response scenarios. Such event-driven logic minimizes the delay between a change in the host's security context and the subsequent restriction of its access.

Reducing the risk of infrastructure compromise often requires additional restrictions on access for legitimate users. These restrictions may include permitting connections only from company-managed devices using machine authentication and validating their security posture. Device posture verification combined with user multi-factor authentication establishes a highly secure access standard for corporate infrastructure. Situations requiring access from non-corporate

devices necessitate special identification and differentiation measures. Contextual policy technologies make it possible to design various identification scenarios for partner devices.

For example, Host Information Profile (HIP) technology embedded in VPN clients enables verification of operating system parameters, security settings, and installed applications. In such cases, access control may rely on an agreed set of mandatory applications that identify a device as a partner device, verification of a unique registry key, or validation of a device serial number.

An alternative approach involves the use of third-party applications that assess device posture. For instance, an antivirus agent or a mobile device management system can provide device state information to a SOAR platform, which—using REST APIs—registers IP addresses in a dynamic firewall address group with a specific tag such as “partner_device” or “IOT_corporate”. This group is then referenced by preconfigured firewall policies that either allow or block access accordingly.

Dynamic address group technologies enable host addresses to be registered into predefined firewall rules and allow access changes to occur automatically as a result of detected events or changes in device state. This capability supports the automation of response processes to evolving cybersecurity conditions, thereby significantly increasing the overall resilience of the infrastructure against cyber threats.

An important and valuable capability of next-generation firewalls is their ability to recognize application-layer (Layer 7) fingerprints within network traffic and to enforce access control not only based on service–port combinations but also based on applications themselves. Application awareness significantly enhances access control flexibility and improves the effectiveness of network access management. Technologies that identify applications within IP packets allow more precise and effective control over access to different applications and services, even when they use the same ports and protocols for data transmission.

The following examples illustrate scenarios where application-aware controls provide value both in cybersecurity and in improving the flexibility and efficiency of network access management:

- visibility into traffic tunneling applications helps minimize potential scenarios of uncontrolled data exfiltration through unmanaged VPN tunnels;
- allowing recreational applications according to a schedule can improve workforce productivity during business hours;
- application-based access rules enable access to services with dynamic ports or addresses without explicitly enumerating them, preserving the principle of least privilege even in highly dynamic environments.

4. Results

The conducted research has shown that the use of contextual policies in organizing network access control within a Zero Trust Architecture plays a key role in ensuring compliance with the fundamental principles of Zero Trust. The application of user account recognition technologies during connection requests enables network access control based on identity, thereby extending beyond the limitations of IP address–based access control.

The use of user or user group objects in security policies makes it possible to reduce the number of security policies by at least half, taking into account a minimum of two user network entry points- internal and remote access. This approach also enables more rigorous enforcement of the principle of least privilege compared to policies based solely on IP addresses and ports. The most significant contextual attributes include identity verification, device posture validation, and application-layer (Layer 7) recognition and firewalling by next-generation firewalls, as illustrated in figure 4.

The study further identified the optimal depth of access level organization for a single network node when applying contextual verification mechanisms as no more than three levels, considering configuration complexity and troubleshooting overhead. In addition, the research confirmed the

inefficiency of the practice of defining network access exceptions layered on top of broadly defined security rules.

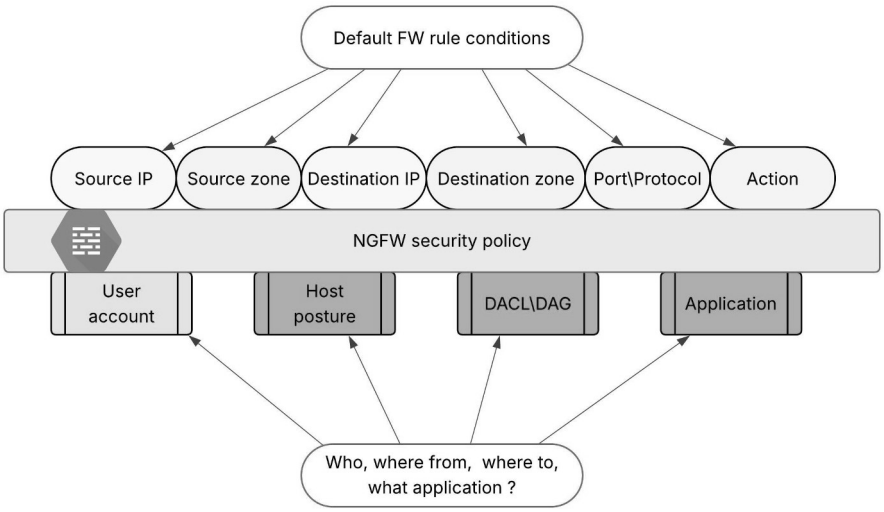

Figure 4. Standard and contextual attributes of access policies

The use of contextual validation in combination with dynamic access lists and dynamic address group technologies has demonstrated the potential to build scenarios for the automatic adjustment of a host's access level in response to changes in its security posture.

According to the results of the study on integrating mechanisms for collecting and enforcing contextual attributes, despite the security benefits of applying such mechanisms within firewall security policies, they also introduce certain drawbacks. These include increased total cost of ownership and operational maintenance of the information infrastructure, potential degradation of network performance due to dependencies on external systems, and the introduction of risks associated with moving access control decision points outside the corporate firewall.

5. Conclusions

The significance of network connection context in access policy management manifests not only in the security domain but also in operational efficiency. Neglecting contextual factors typically leads to difficulties in regulating access within dynamic environments, resulting in the accumulation of rules that are difficult to maintain and that introduce potential vulnerabilities. In contrast, the implementation of additional contextual validations enables more precise definition of required access and reduces the number of redundant rules. A key requirement is the continuous verification of identity, device posture, and network environment attributes, which enables real-time access decision-making. As a result, contextual policies extend beyond traditional approaches and create a more granular, flexible, and adaptive access control system.

Consequently, both the level of protection of critical information resources and the operational efficiency of administrators are improved, making contextual policies one of the key instruments for implementing Zero Trust Architecture.

Among the drawbacks of the proposed advanced access control methods is the increased complexity of the infrastructure. Dependencies on external systems may affect firewall performance, while insufficient protection of context collection and delivery mechanisms expands the attack surface and increases the risk of unauthorized influence and compromise of both network and security infrastructures. Additionally, rule design methodologies that support multi-level access control or the practice of defining contextual exceptions to baseline policies can complicate operational support and increase the effort required for access troubleshooting.

Future research directions include the analysis of risks associated with unstructured network access management and the development of a comprehensive lifecycle concept for managing network access in corporate infrastructures.

Declaration on Generative AI

While preparing this work, the authors used the AI programs Grammarly Pro to correct text grammar and Strike Plagiarism to search for possible plagiarism. After using this tool, the authors reviewed and edited the content as needed and took full responsibility for the publication's content.

References

[1] C. Ike, et al. Redefining zero trust architecture in cloud networks: A conceptual shift towards granular, dynamic access control and policy enforcement. Magna Scientia Advanced Research and Reviews, 2021, doi:10.30574/msarr.2021.2.1.0032

[2] K. Routray, P. Bera. ZTAAC: Zero trust adaptive authorization with CP-ABE for context-aware data protection, in: 2025 17th Int. Conf. on Communication Systems and Networks (COMSNETS), 2025, 814–816, doi:10.1109/COMSNETS63942.2025.10885763

[3] S. Ahmadi. Autonomous identity-based threat segmentation in zero trust architectures. ArXiv, 2025, abs/2501.06281.

[4] D. Yaganti. Securing .Net microservices through conditional access and zero trust principles using Azure AD and OAUTH2. Int. J. of Advanced Research in Science, Communication and Technology, 2023, doi:10.48175/ijarsct-18000a

[5] M. Hirai, D. Kotani, Y. Okabe. Linking Contexts from Distinct Data Sources in Zero Trust Federation, 2022, 136–144. doi:10.48550/arXiv.2209.11108

[6] O. Oluoha, et al. AI-enabled framework for zero trust architecture and continuous access governance in security-sensitive organizations. Int. J. of Social Science Exceptional Research, 2024, doi:10.54660/ijsser.2024.3.1.343-364

[7] S. Xiao, et al. SoK: Context and risk aware access control for zero trust systems. Security and Communication Networks, 2022. doi:10.1155/2022/7026779

[8] O. Ejiofor, O. Olusoga, A. Akinsola. Zero trust architecture: A paradigm shift in network security. Computer Science & IT Research Journal, 2025, doi:10.51594/csitrj.v6i3.1871

[9] O. Harasymchuk, et al. Intelligent cyber defence systems: detection of ransomware and protection of wireless networks based on artificial intelligence technologies: collective monograph. Technology Center PC, 2025. doi:10.15587/978-617-8360-22-1

[10] N. Stojanovski, M. Gusev. Architecture of a identity based firewall system. ArXiv, 2011, abs/1108.1344.

[11] T. Dimitrakos, et al., Trust aware continuous authorization for zero trust in consumer Internet of Things, in: 2020 IEEE 19th Int. Conf. on Trust, Security and Privacy in Computing and Communications (TrustCom), Guangzhou, China, 2020, 1801–1812, doi:10.1109/TrustCom50675.2020.00247

[12] A. Fadhel, D. Bianculli, L. Briand, B. Hourte. A model-driven approach to representing and checking RBAC contextual policies, in: 6th ACM Conf. on Data and Application Security and Privacy, 2016, doi:10.1145/2857705.2857709

[13] S. Panda, S. Sahoo, R. Halder, S. Mondal. Contextual attribute-based access control scheme for cloud storage using blockchain technology. Software: Practice and Experience, 54, 2023, 2042–2062, doi:10.1002/spe.3250

[14] O. Sapsai, Y. Martseniuk, A. Partyka, O. Harasymchuk. Research on automated security incident management in public cloud environments, in: CEUR Workshop Proceedings, 4042, 226–249.

[15] V. Chornii, Y. Martseniuk, A. Partyka, O. Harasymchuk. Information security risks associated with the uncontrolled storage of secrets in source code, in: CEUR Workshop Proceedings, 4042, 250–271.
[16] P. Skladannyi, et al., Improving the security policy of the distance learning system based on the zero trust concept, in: Cybersecurity Providing in Information and Telecommunication Systems, 3421, 2023, 97–106.
[17] Y. Martseniuk, A. Partyka, I. Opirskyy, O. Harasymchuk. Cost observability as a security control In Multi-cloud enviroments based on SOC 2 Security standart. Computers & Security, 2025, 104771. doi:10.1016/j.cose.2025.104771