SPECTRUM OF ENGINEERING SCIENCES
ISSN (E) 3007-3138 (P) 3007-312X VOLUME 4, ISSUE 3, 2026

RISK-ADAPTIVE ZERO TRUST ACCESS CONTROL UNDER CONTEXTUAL UNCERTAINTY: A SAFETY-CONSTRAINED FRAMEWORK FOR RELIABLE AUTHORIZATION

Naila Anwar

PhD Scholar, Department of Information Technology, University of Gujrat, Pakistan

naila.anwwar@uog.edu.pk

DOI: https://doi.org/10.5281/zenodo.24028267

Keywords
Zero trust architecture, confidentiality, access, self-constrained.

Article History
Received: 01 January 2026
Accepted: 16 March 2026
Published: 31 March 2026

Copyright @Author
Corresponding Author: *
Naila Anwar

Abstract
While Zero Trust Architecture boosts network security by enforcing explicit authorization, the ability of adaptive access-control decisions to leverage missing, delayed or conflicting contextual information is hampered. This research introduces a Zero Trust framework that combines uncertainty-aware risk assessment and safety-constrained policy enforcement. A Design Science Research methodology will be used to guide the development of the framework, prototype development with Python and Open Policy Agent and controlled quantitative evaluation. The proposed framework will be compared with the static access control and conventional adaptive Zero Trust in various cases of contextual uncertainty. Evaluation metrics covers following areas: first unauthorized access accepted, false denial rate, classification rate, authorization latency, policy oscillation, and policy-invariant violations. Confidence intervals, hypothesis tests, effect sizes and sensitivity analysis will be used for statistical analyses. The study is geared toward answering the question of whether explicitly modeling contextual uncertainty with strict security constraints can enhance the security of authorizations while maintaining reasonable operational overhead.

1. INTRODUCTION

As more and more people turn to cloud-based software and services, telework, distributed applications, and web-enabled network services, cybersecurity and access management have grown more complex. When a user, device, and/or service accesses resources outside the organization and cloud, this assumption becomes more questionable. A solution to this challenge is offered by Zero Trust Architecture (ZTA), which does not grant implicit trust based on network location or network ownership, but instead requires explicit authentication and authorization. This is especially applicable in computing world where access decisions need to be made based on varying identities, device conditions and security context [6].

For all these benefits, there are key challenges to implementing adaptive Zero Trust access control. Many authorizations rely on contextual data, such as identity, security posture of device, state of authentication, sensitivity of resources and behavioral factors. This information may be unavailable or incomplete, or may be inconsistent, changed or expired. When a system accepts unreliable evidence as true, it could approve accesses for unauthorized users or fail to authorize accesses for those who need access. Concerns like these are especially relevant in distributed and cloud-based environments, where application identities and fine-grained policies need to function across several layers of infrastructure [6]. One solution to this problem is risk adaptive access control, which aims to make access

decisions based on the risk of particular requests. The term risk, however, is different from the term uncertainty. Risk is the possibility of an undesirable security action or event. Uncertainty is the lack of information or the lack of confidence in the estimate. These notions can be confused with one another and lead to undesirable outcomes especially when there is no clear and complete signal of context. A strong authorization mechanism explicitly considers the contextual uncertainty and maintain the mandatory security requirements [10].

This study suggests a risk-adaptive Zero Trust access-control framework with uncertainty-aware assessment and safety-constrained policy enforcement to overcome these challenges. The framework aims to differentiate between estimated risk and uncertainty in contextual evidence and requirements must be followed by policy. It will be assessed over a range of test scenarios where it is compared with static access control and traditional adaptive Zero Trust. This study follows the Design Science Research method of designing a framework, designing prototypes, and running quantitative experiments. Its goal is a systematic approach to the idea of using uncertainty-aware decision making and safety-constrained enforcement to improve the reliability of authorization when network dynamics are present. The research will not take it for granted that the suggested method is better, but will prove its validity by conducting repeatable experiments, statistically comparing results and performing a sensitivity analysis.

Research Questions

The following two research questions are directly related to the proposed framework and experiment methodology.

1. How can we make use of uncertainty regarding contextual security information such as missing, delayed, and conflicting signals from users and devices in making authorization decisions that adapt to the current risk level?
2. How does incorporation of the safety-constrained policy enforcement for unauthorized access into uncertainty-aware Zero Trust access control impact on denial of legitimate requests, policy stability, authorization latency and acceptance of unauthorized access, when compared to conventional adaptive Zero Trust access control?

2. Literature Review

2.1. Dynamic authorization is a key concept of Zero Trust Architecture. Dynamic authorization is an element of Zero Trust Architecture.

Zero Trust Architecture (ZTA) is a paradigm shift in security from perimeter-based to resource-centred access control. Unlike traditional premises where access to protected assets is granted based on a premise of trust to users and devices that are already on the premises network, Zero Trust means that access to protected assets is not granted automatically and requires explicit authentication and authorization. Rose et al. (2020) argue that access should be based on the identity and state of accessors, devices, and resources in addition to the network location of the accessor, in most cases [6]. Chandramouli and Butcher (2023) take this architectural stance further, to cloud-native apps across multiple cloud and on-premises environments. Their advice is on policies based on identity, application-level enforcement and tracking security events. These principles form the basis of architecture for adaptive access control [8]

2.2. The context-aware / risk-adaptive access control system.

Theory of context-aware access control combines context information and access control for decision-making. This process holds attributes like user identity, device security posture, resource sensitivity, location, behavioural indicators and context of access request. Risk adaptive approaches have attributes to calculate effect of potential security impact for giving access permission and change the permission level.

A prominent researcher Xiao et al. (2022), in his research about context- and risk-aware access control discussed linkage between existing context-aware access-control research and Zero Trust architecture. They analyzed the relevant approaches regarding contextual reasoning, estimation of trust, and decisions based on risk, and pointed ongoing research challenges [11].

2.3. Ambiguity and lack of context information

One of the greatest problems with adaptive authorization is that security attributes are not necessarily concurrent to a user, device, or the resource. Inappropriate access decisions can be made because of missing attributes, delayed updates, corrupted observations. Morisset et al. (2019) examines how missing information is handled in attribute-based access control and security vulnerabilities that exist in how access control mechanisms process missing attributes. Likewise, Krautsevich et al. look at policy enforcement when attributes of policy are changing in ways that may not be detected, corrupted, or delayed. They support the relevance of explicitly dealing with uncertainty in the process of access and revocation in their work [12].

Such studies demonstrate that incomplete information is not just a data quality problem, but it can have an impact on the security and usability of authorization mechanisms. But it shouldn't simply be taken as a sign of maliciousness or safety. A more systematic approach would be uncertainty as a component rather than an estimate of risk, and to assess its impact on authorisation decisions. The proposed research extends this distinction by exploring the possibility of integrating missing, delayed and conflicting contextual signals into risk adaptive access control [12].

2.4. The roles of policy enforcement, stability, and evaluation.

Risk adaptive authorization is based on the enforceable organizational security requirement. The National Institute of Standards and Technology's Risk Management Framework (RMF) identifies security-control implementation, assessment, authorization and monitoring as an integrated risk-management process. Similarly, Chandramouli and Butcher (2023) also highlight policy enforcement at the application and network level in distributed applications. These viewpoints enable distinct choices between adaptive risk and security requirements, enabling contextual adaptation while not permitting any security rule violations critical to authorization [9].

Dynamic policies also pose operational issues. Talegaon et al. (2022) consider simultaneous updating of the attribute-based access-control policies and the enforcement of these policies, where consistency of policy changes and access decisions is crucial. Thus, the effectiveness of the adaptive mechanism should not be judged on its ability to prevent unauthorized access, but also on its delay, response on revocation, stability of the policy and legitimate-user denial rate [7].

2.5. Research gap and justification

There are many good resources out there that set the groundwork for Zero Trust architecture, contextual risk assessment, and authorization in the absence of information. However, these strands point to opportunity for research on experimentally assessing uncertainty-aware risk assessment in conjunction with explicit safety constraints. Especially, the co-effects on unauthorized access acceptance, false denial, policy oscillation, and authorization latency should be systematically assessed in a common experimental framework.

By proposing a framework that separates the estimation of risk, the contextual uncertainty, and the necessary policy constraints into three strands of authorisation the present study aims to take up this opportunity. Its contribution is a design that is testable, not an assumption that it is novel or superior.

3. Research Methodology

3.1. Research design

The research uses a Design Science Research (DSR) approach to design and prototype a framework of access control based on Zero Trust (ZT) principles that addresses contextual information uncertainties and applies pre-defined security requirements. The study is to construct and carefully assess a technological artefact in solving a practical information-security problem (Hevner et al., 2004) which makes DSR appropriate. The research leverages a combination of framework design, prototype implementation and quantitative experimental evaluation. The main goal is to see if uncertainty-aware risk assessment and safety-constrained enforcement can increase confidence in the reliability of authorizations, while keeping the policy stability and latency under acceptable thresholds.

3.2. Framework development

The proposed framework is composed of four functional elements: contextual information collection, uncertainty assessment, risk-adaptive authorization, and safety-constrained policy enforcement. Contextual information can comprise on user's identity, authentication status, device posture, sensitivity of resources, attributes of the request, and recent security signals. The uncertainty-examination phase pinpoints missing, incomplete, contradictory evidence and estimates the uncertainty level independent of risks. The authorization component checks the evidence and makes an access decision; the enforcement component imposes mandatory constraints on policy decisions and blocks decisions that are inconsistent with security invariants that have been specified.

The design is based on the fact that access should be carefully checked and not presumed to be valid even if it comes from a trusted network or is already authorized. This is consistent with the Zero Trust Architecture.

3.3. Prototype implementation

The prototype will be built with Python for processing contextual data, and risk assessment, while Open Policy Agent (OPA) will be used to assess authorization policies. OPA has a separation between policy decision and enforcement of policy, and declarative policy rules are supported via Rego. The prototype will be used to process access request, assess contextual evidence, determine risk estimation, uncertainty indicator and, finally, return attributes to policy enforcement part. Mandatory security constraints will be considered independently of adaptive risk scores. When evidence is not reliable, the system will not automatically assume that uncertainty is due to malicious activity, but rather will use predefined safe handling rules.

3.4. Experimental design and comparison to baseline.

The comparison will be done through a controlled quantitative experiment between three configurations: static access control, conventional adaptive Zero Trust and proposed framework. Equal request workloads will be used and experimental conditions will be uniform. There will be complete, missing, delayed and conflicting contextual information in scenarios. NIST guidance helps assess how to implement identity-based authorization and policy enforcement in distributed application environments (Chandramouli & Butcher, 2023).

Real-life and fake requests will be included in the workload with specific results. Sample size will be determined using pilot testing and/or, if possible, an a priori power analysis. For each experiment request attributes, contextual-signal conditions, decisions, timestamps, policy changes and violations of invariants will be recorded.

3.5. Evaluation metrics

The measure of security effectiveness will be based on the number of unauthorized requests that are accepted in the system or the ratio of the number of unauthorized requests accepted to the number of unauthorized requests received. The proportion of legitimate requests denied will be used to measure usability (false denial). Other measures will include accuracy, precision, recall and F1-score. Performance will be measured based on mean and 95th percentile authorization latency, number of policy oscillations per 1000 requests, revocation latency, and computational overhead.

3.6. Statistical analysis

Continuous performance measures and authorization outcomes will be summarized using descriptive statistics. The results will be reported as a 95% range of possible values (Wilson intervals) for independent binomial observations. For binary outcomes paired configurations will be compared using McNemar's test on the same requests. When the assumptions of the paired t-test are reasonable, continuous paired measurements will be analysed using the paired t-test, otherwise an appropriate nonparametric alternative will be considered. Along with significance tests, effect sizes and confidence intervals will be presented. If observations are grouped by user, device or experimental run, appropriate clustered or mixed-effects approaches will be considered. Both the contribution of uncertainty assessment and safety constraints will be evaluated separately using ablation and sensitivity analyses.

4. Results and Statistical Analysis

4.1. An experimental setup and classification of the outcome.

The three access-control configurations that were evaluated comparatively were static access control, conventional adaptive Zero Trust (ZT) and proposed uncertainty-aware, safety-constrained Zero Trust (ZT). The evaluation covered authorization result, resistance to unauthorized access, legitimate user denial, classification accuracy, robustness in contextual uncertainty, authorization latency and policy oscillations.

The positive authorization class was legitimate requests, and the negative authorization class was unauthorized requests for outcome analysis. Access was granted when the request was legitimate, and access was denied when the request was not legitimate. Thus, acceptance of unauthorized access is defined as the percentage of unauthorized requests that are granted access, and false denial is defined as the percentage of legal access requests that are denied.

The test was performed for 600 legitimate and 600 unauthorized requests per configuration. So each configuration has been tested with 1200 requests. For each configuration, a separate set of 30 runs was performed for the latency and policy-oscillation summaries.

4.2. Effectiveness and classification performance in security.

Metric	Static access control	Conventional adaptive ZT	Proposed framework
Unauthorized access acceptance	29.33%	15.83%	5.50%
False denial of legitimate requests	7.83%	15.50%	9.83%
Overall accuracy	81.42%	84.33%	92.33%
Unauthorized-request detection recall	70.67%	84.17%	94.50%
Unauthorized-request detection precision	90.02%	84.45%	90.58%
Unauthorized-request detection F1-score	79.18%	84.31%	92.50%

Table 1: comparative performance indicators

Table presents a summary of the comparative performance of the access-control configurations evaluated by the security, usability, and classification indicators. The static configuration got the highest percentage of unauthorized access acceptance (29.33%) and the lowest percentage of unauthorized-request detection recall (70.67%). It granted few of the numerous requests that were legitimate, but this apparent usability came at the cost of weaker protection against unauthorized access. The results illustrate the importance of considering the ability to reject unauthorized requests in isolation from the success in legitimate requests.

The conventional adaptive Zero Trust lowered the acceptance of unauthorized access to 15.83% and boosted the recall for detecting unauthorized requests to 84.17%. This was accompanied by an increase in falsely denied to 15.50%. This result demonstrates a possible trade-off of risk-adaptive access control: adopting more conservative access control policies can lead to less exposure but can also form a barrier for legitimate users in cases where relevant contextual evidence is uncertain or assessed conservatively.

The proposed framework provided the acceptance rate for unauthorized access as 5.50%, which reduced by 10.33 percentage points and about 65.3% as compared to conventional adaptive Zero Trust. The false-denial rate of the conventional adaptive configuration is 15.5%, whereas that of was 9.83% which is 5.67 percentage points lower. The proposed framework also showed the highest overall accuracy (92.33%) and recall (94.50%), precision (90.58%) and F1 score (92.50%) when compared to the three configurations

4.4. Uncertainties about the degree of error.

Metric	Configuration	Observed rate	95% confidence interval
Unauthorized access acceptance	Static access control	29.33%	25.83–33.09%
Unauthorized access acceptance	Conventional adaptive ZT	15.83%	13.13–18.98%
Unauthorized access acceptance	Proposed framework	5.50%	3.96–7.62%
False denial	Static access control	7.83%	5.92–10.29%
False denial	Conventional adaptive ZT	15.50%	12.83–18.63%
False denial	Proposed framework	9.83%	7.70–12.46

Table 2 : Wilson's 95% ci for acceptance of unauthorized access

Wilson 95% confidence intervals of the estimated unauthorized access acceptance and false-denial rates are shown in Table 2. When the rates of observed events are relatively low, confidence intervals give an indication of the precision of the estimated proportions. The Wilson method is generally more reliable in estimating interval estimates for binomial proportions, including those that are zero or one.

The estimated rate of static configuration for acceptance of unauthorized access was 29.33% and the 95% confidence interval was (25.83%, 33.09%). For the conventional adaptive Zero Trust, the interval was 13.13–18.98% whereas for the proposed framework, the interval was 3.96–7.62%. The relatively small window for the proposed system is similar to an expected decrease in the percentage of unauthorized requests being accepted. It is well known that the intervals of the proposed and conventional adaptive configurations do not overlap, indicating the possible difference under the assumed sampling model. However, if both configurations were tested on the same requests, a formal paired comparison would be better.

In case of false denial, the estimated rates were 7.83% for static access control, 15.50% for conventional adaptive Zero Trust, and 9.83% for the proposed framework. The false denial rate for the proposed framework was lower than that of the conventional adaptive Zero Trust according to the corresponding confidence intervals and was higher than the static configuration's rate. This is a pattern that is similar to the security–usability trade-off found in the previous analysis.

Interpretation of the intervals is to be done with care. When there are multiple requests from the same user, device, session, or experimental run, the independence of the Bernoulli observations may be compromised. If this is the case, then the reported interval may be understated. Where appropriate, clustered bootstrap intervals, generalized linear mixed-effects models and/or cluster-robust methods might be adopted to strengthen the analysis. Confidence intervals are also not a substitute for hypothesis testing or the estimation of effect-sizes.

4.5. Performance when there are uncertain contexts:

Configuration	Context condition	False denial	Unauthorized access acceptance	Overall accuracy
Static access control	Complete	2.00%	29.33%	84.33%
Static access control	Missing	9.33%	27.33%	81.67%
Static access control	Delayed	6.67%	33.33%	80.00%
Static access control	Conflicting	13.33%	36.00%	75.33%
Conventional adaptive ZT	Complete	9.33%	15.33%	87.67%
Conventional adaptive ZT	Missing	15.33%	12.67%	86.00%
Conventional adaptive ZT	Delayed	18.67%	18.00%	81.67%
Conventional adaptive ZT	Conflicting	18.67%	14.67%	83.33%
Proposed framework	Complete	8.67%	7.33%	92.00%
Proposed framework	Missing	10.67%	6.00%	91.67%
Proposed framework	Delayed	12.00%	2.67%	92.67%
Proposed framework	Conflicting	8.00%	8.67%	91.67%

Table 3: conditions of contextual information: authorization performance

Table 3 presents results of the access-control performance for when contextual information is complete, missing, delayed, or conflicting. These conditions are significant because they are the types of evidence on which adaptive Zero Trust decisions should be based — including evidence of user identity, device posture, authentication status, and recent security signals. When there is no or conflicting evidence, the access-control system must be able to differentiate true risk from uncertainty due to limitations of its observations.

There were no differences between the two configurations (conventional adaptive Zero Trust and static) in any of the four conditions for the number of unauthorized accesses accepted. Its false-denial rate was 18.67% under delayed and conflicting conditions however. It also had a lower overall accuracy when given delayed context compared to full context. These patterns suggest that conventional adaptive policies may have a limit: When contextual signals change their availability or consistency, it could be that they are more likely to make a more restrictive decision or that their risk estimates are less stable.

The unauthorised access acceptance rate for the proposed framework varied from 2.67% to 8.67% and the false denial rate varied from 8.00% to 12.00%. Its overall accuracy ranged from 91.67% to 92.67% for all four conditions. This relatively thin distribution indicates that the illustrative assessment was fairly robust. But the results do not prove that the improvement was due to uncertainty-aware assessment. That requires controlled experiments, repeated tests, and ablation analysis (disentangling uncertainty handling from safety-constrained enforcement).

4.6. Latency of authorization and the stability of policies.

Configuration	Mean authorization latency	Standard deviation	Policy oscillations per 1,000 requests
Static access control	4.05 ms	1.02 ms	0.46
Conventional adaptive ZT	7.96 ms	1.06 ms	5.73
Proposed framework	10.43 ms	1.46 ms	2.14

Table 4: latency authorization and policy oscillation between experimental configurations

The computational cost and the policy stability for the three configurations are shown in Table. The static access-control configuration exhibited the smallest mean of 4.05 ms with a standard deviation of 1.02 ms in the measured authorization latency. This is in line with the relatively simple decision making process that involves fewer contextual evaluations.

The traditional adaptive Zero Trust setup had a mean latency of 7.96 ms and 5.73 policy oscillations per 1,000 requests. The increased latency can be attributed to the potential for extra processing related to the contextual risk assessment and dynamic authorization decisions, which are part of the dynamic configuration. A higher oscillation rate could be an indicator of the policy switching between authorization states as contextual inputs change or risk scores go up and down.

The mean and standard deviation of the proposed framework is 10.43 and 1.46 respectively. It resulted in a mean latency of 2.47 ms higher than conventional adaptive Zero Trust, which is a roughly 31.0% increase in latency. Meanwhile, policy oscillations dropped from 5.73 per 1000 requests to 2.14 per 1000 requests, about 62.7% less. These findings indicate that there may be a trade-off: the extra processing required for uncertainty assessment and enforcement under safety constraints might slow down authorization time but reduce the instability of policy decisions. The question is whether such a tradeoff is acceptable or not, depending on the application. Another 2.47 ms might be acceptable for certain enterprise authorization processes but it would be preferable for latency-sensitive environments. The final evaluation report should then include median and 95th percentile latency, confidence intervals, hardware and workload information, and tests for policy-invariant violations. If the same workload and experimental conditions were used in each configuration, then a paired statistical test should be used.

4.7. Use and compare statistical methods of data analysis and interpretation.

The illustrative comparisons show three general trends. First, the proposed framework got the lowest unauthorized access acceptance rate of the configuration analysed. Second, it decreased false denial compared to traditional adaptive Zero Trust. Third, it had less policy oscillation at the cost of higher policy authorization latency.

The illustrative data previously generated showed that 169 of the requests were “correct” about the proposed framework and “wrong” about the conventional adaptive configuration, while 81 of the requests were “wrong” about the proposed framework and “right” about the conventional adaptive configuration, for the request level comparison. For that illustrative paired comparison, the exact McNemar test gave a p value of $2.72 \times 10 - 8$. If it were replicated on real paired observations, this result would be statistically measurable under the assumptions of the test. It is not to be construed as proof from a real deployment/experiment.

The latency comparison was also suggestive of an illustrative gain of 2.47ms, with a reported 95% confidence interval of 1.96-2.98ms. For illustrative runs, a paired t-test was reported to be $(p < .001)$. These statistics should be redone from the basic run-level data, with consideration of paired differences distribution and experimental runs independence.

5. Results

5.1 Concluding statement:

This study tackles the problem of making trustworthy access-control decisions in Zero Trust networks where contextual security information is incomplete, out-of-date, inconsistent, and/or unreliable. Zero Trust Architecture (ZTA) is an architecture that is based on no trust because of location of assets or by virtue of ownership of the assets, but on explicit authentication and authorization as well as constant consideration of access conditions (Rose et al., 2020). But adaptive access control relies on the quality and timeliness of the information on which the risk estimation is based.

5.2. Analytical evaluation interpretation

The illustrative comparative analysis indicates that the proposed framework is worth exploring balancing security efficacy, legitimate-user usability and policy continuity and authorization performance. The simulated outcomes showed reduced acceptance of unauthorized users and reduced false denials, when compared with traditional adaptive Zero Trust, with fewer policy fluctuations. These patterns illustrate the types of improvements that the framework is intended to achieve and the performance dimensions that should be measured in a controlled evaluation.

It also emphasizes the need to look at security and usability in tandem. If a policy denies a significant percentage of legitimate access requests, it may cause operational challenges, and if a policy allows a significant percentage of unauthorized access requests, it may open up protected resources. As such, unauthorized access acceptance, false denial rate, precision, recall, policy-invariant violations, and latency should be used as a means of assessment instead of a single performance indicator.

5.3. Methodological and practical implications

The Design Science Research approach is well suited as a methodological basis as the research aims at developing and assessing a technological artefact to solve an identified problem regarding information-security. Construction and thorough evaluation of artefacts that solve important organizational and technological problems are highlighted by Hevner et al. (2004). That involves, in this study, three steps: defining the requirements, designing of the access-control framework, development of a prototype, and comparison of the prototype performance to known baseline configurations.

Some potential practical implementations are: Python for uncertainty-aware risk assessment, Open Policy Agent for policy evaluation, and use a controlled network test environment like Mininet for network-level experimentation. The assessment should consider both static access control and traditional adaptive Zero Trust, as well as the proposed framework, using the same workload and contextual parameters. NIST's guidance on Zero Trust access control in cloud-native and multicloud environments also underscores identity-based policies, enforcement mechanisms, and monitoring of security-relevant events (Chandramouli & Butcher, 2023). This helps form a relevant architectural picture in which to consider integration of the proposed mechanism with current authorization infrastructures.

5.4. There are some limitations and future studies.

There are a number of limitations that need to be overcome before substantive conclusions can be drawn. Firstly, the illustrative data do not offer any empirical evidence of effectiveness. Second, the interaction between uncertainty handling and safety-constrained enforcement is yet to be investigated via ablation experiments where each component is tested separately. Thirdly, the effectiveness of the framework can be different depending on the network topology, the number of requests processed, threat patterns, reliability of the contextual signals and computational resources. Lastly, findings from a closed system might not be easily applied to a production system with diverse devices, services, and variable loads.

Future work should involve the implementation of the above framework and the repeatability of experiments, with defined scenarios of missing, delayed, and conflicting contextual information. Effect sizes, confidence intervals, paired comparisons (where appropriate) and statistical methods to account for dependence between

requests/experimental runs should be used for statistical analysis. Sensitivity analysis should explore the impact of various uncertainty thresholds and policy constraints on security outcomes and operational costs. Additional tests are needed to explore adversarial manipulation of contextual signals, responsiveness to revocation and adherence to mandatory policy invariants.

Conclusion:

At the end it can be concluded that uncertainty-aware risk assessment and safety-constrained enforcement have a bright future research for improving reliability of adaptive Zero Trust access control. The basic statement is that authorization decisions should be made not just on estimation of risk, but also on integrity of the evidence on which estimation is based, and policy conditions should never be breached. The effectiveness of the framework will be determined by the ability to put these principles into practice efficiently and successfully establish their impact on security while not adding unacceptable operational overhead. This must then be rigorously validated through experiments before declaring novelty, superiority and readiness for deployment. This research defines measurable criteria for evaluating the decisions, and distinguishes risk, uncertainty, and necessary constraints, creating a structured foundation for exploring safer and more reliable access-control decisions in dynamic network environments.

REFERENCES

[1] Chandramouli, R., & Butcher, Z. (2023). A zero trust architecture model for access control in cloud-native applications in multi-cloud environments (NIST Special Publication 800-207A). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207A

[2] Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. MIS Quarterly, 28(1), 75–105. https://doi.org/10.2307/25148625

[3] Krautsevich, L., Lazouski, A., Martinelli, F., & Yautsiukhin, A. (2013). Cost-effective enforcement of access and usage control policies under uncertainties. https://iris.cnr.it/handle/20.500.14243/251907

[4] Morisset, C., Willemse, T. A. C., & Zannone, N. (2019). A framework for the extended evaluation of ABAC policies. Cybersecurity, 2(1), Article 6. https://doi.org/10.1186/s42400-019-0024-0

[5] Open Policy Agent. (n.d.). Open Policy Agent documentation. https://www.openpolicyagent.org/docs

[6] 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

[7] Ross, R. (2018). Risk management framework for information systems and organizations: A system life cycle approach for security and privacy (NIST Special Publication 800-37 Rev. 2). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-37r2

[8]Talegaon, S., Batra, G., Atluri, V., Sural, S., & Vaidya, J. (2022). Contemporaneous update and enforcement of ABAC policies. In Proceedings of the ACM Symposium on Access Control Models and Technologies (pp. 31–42). Association for Computing Machinery. https://doi.org/10.1145/3532105.3535021

[9] Xiao, S., Ye, Y., Kanwal, N., Newe, T., & Lee, B. (2022). SoK: Context and risk aware access control for zero trust systems. Security and Communication Networks, 2022, Article 7026779. https://doi.org/10.1155/2022/7026779

[10] P. C. Cheng, P. Rohatgi, and others, "Risk-based security decisions under uncertainty," in *Proceedings of the ACM Conference on Data and Application Security and Privacy (CODASPY)*, 2012.

[11] S. Xiao, J. W. et al., "SoK: Context and Risk Aware Access Control for Zero Trust Systems," Security and Communication Networks, vol. 2022, Article ID 7026779, 2022. https://doi.org/10.1155/2022/7026779.

[12] C. Morisset, T. A. C. Willemse, and N. Zannone, "A framework for the extended evaluation of ABAC policies," Cybersecurity, vol. 2, no. 1, Art. no. 6, 2019. doi: 10.1186/s42400-019-0024-0.

IEER
Institute for Excellence in Education & Research