

Date: 5th October-2026

IMPROVING ACCESS CONTROL MODELS FOR CONFIDENTIAL DOCUMENTS IN MILITARY INFORMATION SYSTEMS

Giyosjon Jumaev Abdivaxobovich¹, Turapov Sherzod Nurmatovich²

¹Artificial intelligence and digital technologies, Tashkent International University
of Financial Management and Technologies,

²Head of Department of Act and Military Communication Institute
Tashkent, Uzbekistan

¹giyosjonjumaev@gmail.com

Abstract: Background. Classified document systems in defence organisations still rely largely on lattice-based mandatory access control (MAC) and role-based access control (RBAC). These models decide on static labels and ignore need-to-know, the state of the requesting device, and behavioural risk, which leaves them exposed to credential compromise, insider misuse and bulk aggregation. **Objective.** To design and evaluate a hybrid access-control model for confidential documents that combines a Bell–LaPadula lattice, attribute-based access control (ABAC), context and risk evaluation, and a two-person rule for the highest classification. **Methods.** The model was specified formally and compared with MAC, RBAC and static ABAC in a Monte Carlo simulation of 120,000 requests per run (10 runs) over a synthetic population of 1,500 subjects and 6,000 documents, with four attack classes. **Results.** In this synthetic setting the hybrid model granted 0.94% of malicious requests, against 75.11% for MAC, 81.83% for RBAC and 39.93% for static ABAC, while denying 2.59% of legitimate requests. **Conclusions.** Adding contextual and risk-adaptive conditions to a lattice-plus-ABAC core removes most residual attack paths at a modest availability cost. Because the evaluation is simulated, the figures show the direction and approximate size of the effect and must be confirmed with operational data.

Keywords: access control; attribute-based access control; Bell–LaPadula; zero trust; classified information; military information systems; risk-adaptive authorization.

1. Introduction

Military information systems hold documents whose disclosure can cost lives, compromise operations or damage alliances. The access-control foundations for such systems were laid in the 1970s. The Bell–LaPadula model [1] formalised the "no read up, no write down" rules over a lattice of classification levels and compartments [3], and the Biba model [2] added the dual concern of integrity. These models remain the conceptual basis of mandatory access control in national security systems, and they give strong, provable guarantees about information flow between labels.

Operational reality has moved on. Coalition operations, expeditionary networks, mobile devices and contractors mean that the person requesting a document, the device in their hand and the network path are all variable. Three weaknesses of label-only control follow. First, clearance is necessary but not sufficient: a holder of a Secret clearance

Date: 5th October-2026

should still see only the documents their mission requires, which is the need-to-know principle that a pure lattice does not express [13,19]. Second, a stolen credential is indistinguishable from its owner at the label level. Third, many individually authorised reads can add up to a compromise that no single decision reveals, a problem recognised since the Chinese Wall policy [5] and now known as aggregation risk.

RBAC [6,7] improved administration by granting permissions through roles, but roles are coarse and, as Kuhn et al. observed, accumulate attribute-like exceptions over time [20]. ABAC evaluates attributes of the subject, the object, the action and the environment against policy [8]. It generalises DAC, MAC and RBAC [11] and has an extensive research literature [10]. At the same time, zero-trust architecture [9] and the Department of Defense zero-trust strategy [17] argue that every request should be authorised dynamically, using the best available evidence about the requester and the context, rather than by network location or a one-time login.

What has not been shown clearly is how much each additional layer buys. Operators need to know whether adding ABAC to a lattice is enough, or whether context and risk signals are also required, and what each choice costs in denied legitimate work. This article therefore asks: **to what extent does a hybrid model that combines lattice rules, ABAC need-to-know, contextual risk and a two-person rule reduce unauthorised access to confidential documents compared with MAC, RBAC and static ABAC, and at what cost to legitimate users?** The contributions are (i) a formal hybrid model, (ii) a reproducible simulation design with four attack classes, and (iii) a quantitative comparison with a sensitivity analysis of the key risk threshold.

2. Methods

2.1 Study design and threat model

The study is a comparative simulation. Four decision models were implemented and applied to the same stream of access requests, so that differences are attributable to the model and not to the data. The threat model covers four classes of malicious request, chosen because each defeats a different class of control (Table 1).

Table 1. Attack classes used in the evaluation and their share of malicious traffic.

Attack class	Description	Share of malicious	Passes labels?
Insider overreach	Cleared user reads documents outside their missions (no need-to-know)	35%	Yes
Compromised credential	Valid credentials used from an untrusted device, usually from an insecure network	30%	Yes
Privilege escalation	Request for documents above clearance or without required compartments	25%	No
Bulk aggregation	Authorised documents requested at an abnormally high rate	10%	Yes

Date: 5th October-2026

2.2 Comparison with classical models

Table 2 summarises what each model can express. The proposed model keeps the lattice as a non-bypassable floor, which preserves the information-flow guarantees of Bell–LaPadula, and adds conditions that narrow, never widen, what the lattice would allow.

Table 2. Capabilities of the access-control models compared.

Capability	BLP/MAC	RBAC	Static ABAC	Proposed hybrid
Classification and compartments	Yes	Partial	Yes	Yes
Need-to-know (mission)	No	Partial	Yes	Yes
Device and network context	No	No	No	Yes
Behavioural risk and rate	No	No	No	Yes
Step-up authentication	No	No	No	Yes
Two-person rule (Top Secret)	No	No	No	Yes
Formal information-flow guarantee	Yes	No	Not by default	Yes (lattice floor)

2.3 The proposed hybrid model

The architecture follows the policy-enforcement and policy-decision split used in zero-trust designs [9] and is shown in Figure 1. A policy enforcement point (PEP) intercepts every request and asks a policy decision point (PDP), which evaluates the policy using attributes from the attribute stores and signals from the context services. Every decision is written to a hash-chained audit log to support accountability.




Figure 1. Proposed hybrid access-control architecture.

Date: 5th October-2026

Let a subject s request action a on document o . The classification levels are Unclassified < Restricted < Secret < Top Secret. The attributes used are listed in Table 3. A request is permitted only if all of the following hold:

$$\text{Permit}(s, o, a) = L \wedge C \wedge N \wedge (X \vee S) \wedge T$$

where **L** is the Bell–LaPadula simple-security condition, $\text{clr}(s) \geq \text{cls}(o)$; **C** requires the subject's compartments to include all compartments of the document; **N** requires the document's mission to be among the subject's assigned missions; **X** is the context condition, requiring a trusted device, a secure network, a risk score below the threshold τ and a request rate below ρ ; **S** allows a degraded context (untrusted device or network) to be redeemed by successful step-up authentication, provided the document is not Top Secret and the risk and rate conditions still hold; and **T** requires approval by a second authorised person when $\text{cls}(o)$ is Top Secret. The conjunction means the model can only deny more than the lattice does, which preserves the lattice guarantees.

Table 3. Attributes evaluated by the policy decision point.

Category	Attribute	Role in the decision
Subject	Clearance level; compartments; assigned missions	Conditions L, C and N
Object	Classification; compartments; mission tag; handling caveat	Conditions L, C, N and T
Environment	Device posture; network zone	Condition X; degraded values trigger step-up in S
Behaviour	Risk score in $[0,1]$; request rate	Condition X against τ and ρ
Process	Second approver availability	Condition T for Top Secret

2.4 Simulation design

Each run generated 1,500 subjects (clearance distribution 25%, 30%, 30% and 15% from Unclassified to Top Secret; each assigned three of 40 missions; five compartments), six roles with different maximum classification, and 6,000 documents (classification mix 20%, 30%, 35%, 15%). Each run issued 120,000 requests, of which 85% were legitimate and 15% malicious, split among the four attack classes in Table 1. Legitimate requests satisfied the full oracle policy; 7% came from untrusted devices, 5% from insecure networks and 3% were exceptional field use with elevated risk. Step-up authentication succeeded for 92% of legitimate users and 10% of attackers, and a second approver was reachable for 88% of legitimate and 10% of malicious Top Secret requests. Risk scores followed Beta distributions with higher means for compromised-credential and aggregation

Date: 5th October-2026

attacks. The thresholds were $\tau = 0.45$ and $\rho = 0.6$. The experiment was repeated for 10 independent random seeds.

Two outcome measures were used: the false-accept rate (FAR), the share of malicious requests granted, and the false-reject rate (FRR), the share of legitimate requests denied. A sensitivity analysis varied τ from 0.2 to 0.8 (five seeds per value). Results are reported as mean $\pm$ standard deviation over runs. **Important limitation:** the population, the attack mixes and the risk distributions are my own assumptions, not measurements from a real system (see Section 4.3).

3. Results

3.1 Overall performance

Table 4 and Figure 2 show the main result. Label-based MAC admitted three of four malicious requests, because insider overreach, compromised credentials and aggregation all present valid labels. RBAC performed worse on FAR and also rejected legitimate users, since coarse roles both over-grant and under-grant. Static ABAC removed need-to-know violations and halved the FAR relative to MAC, but could not detect misuse of valid attributes. The hybrid model reduced the FAR to below one percent at an FRR of about 2.6%.

Table 4. False-accept and false-reject rates (% , mean $\pm$ SD over 10 runs).

Model	FAR – malicious granted	FRR – legitimate denied
BLP/MAC	75.11 $\pm$ 0.44	0.00 $\pm$ 0.00
RBAC	81.83 $\pm$ 0.46	7.71 $\pm$ 0.52
ABAC (static)	39.93 $\pm$ 0.37	0.00 $\pm$ 0.00
Proposed hybrid	0.94 $\pm$ 0.04	2.59 $\pm$ 0.09




Figure 2. Share of malicious requests granted and legitimate requests denied, by model.

3.2 Results by attack class

The per-class results (Table 5, Figure 3) show why the hybrid design works. Privilege escalation is stopped by every model that includes the lattice. Insider overreach is

Date: 5th October-2026

stopped once need-to-know is expressed, which requires ABAC. Compromised credentials and bulk aggregation pass both lattice and static ABAC in every case and are caught only by context and rate conditions.

Table 5. Share of attacks that succeeded, by attack class (% , mean over 10 runs).

Attack class	BLP/MAC	RBAC	Static ABAC	Proposed hybrid
Insider overreach	100.0	92.3	0.0	0.0
Compromised credential	100.0	92.3	100.0	2.0
Privilege escalation	0.0	50.3	0.0	0.0
Bulk aggregation	100.0	92.1	100.0	3.3




Figure 3. Attack success by class for MAC, static ABAC and the proposed hybrid model.

3.3 Sensitivity to the risk threshold

The risk threshold τ controls the trade-off between security and availability (Figure 4). A strict threshold of 0.2 gave a FAR near zero but denied about 26% of legitimate requests, which would be unworkable in operations. Raising τ to 0.8 cut the FRR to about 2% but the FAR rose to about 4%. The chosen value of 0.45 lies near the knee of the curve, where the FAR is below 1% and the FRR under 3%. Above $\tau \approx 0.5$ the FRR barely improves, because the remaining denials come from the two-person rule and failed step-up rather than from the risk score.




Figure 4. Effect of the risk threshold τ on false-accept and false-reject rates for the hybrid model.

4. Discussion

4.1 Principal findings

Three findings stand out. First, each layer closes a specific gap and none is redundant: the lattice stops escalation, ABAC stops insiders who overreach, and context and risk stop attackers who hold valid credentials. This matches the reasoning behind unified lattice, role and attribute models [11] and behind zero-trust guidance that decisions should use continuous evidence rather than a single login [9,17]. Second, the cost of these protections is modest if the threshold is tuned, with an FRR of about 2.6% in the baseline configuration. Third, RBAC alone performed worst in this setting, which is consistent with the long-standing observation that roles are a poor proxy for need-to-know [20,10].

4.2 Practical implications

For system owners, the results support an incremental path: keep the existing lattice as a floor; add mission and handling attributes to documents and users, which delivers the largest single gain; then add device posture, network zone and a risk score. The two-person rule should be limited to the highest classification, since it is the main source of residual denials and delays. Attribute quality is a precondition for all of this, so attribute management and audit deserve at least as much engineering effort as the policy engine [8,10]. Detailed control families for such deployments are given in NIST SP 800-53 [16], and database-level enforcement should follow the principles in [18].

Residual risks remain even in the model. The hybrid model enforces access, not use after disclosure, so usage-control extensions [12] are relevant for downstream handling. Safety questions about how rights can propagate through a system are undecidable in general [14], which argues for keeping the policy small and auditable. The classical design principles of least privilege, fail-safe defaults and complete mediation [13] still apply to the PEP and PDP themselves, and the foundational notion of a protection matrix [15] remains a useful way to review the resulting policy.

4.3 Limitations

This evaluation has several limitations that readers should weigh heavily. (1) The data are synthetic and the attack mixes, risk distributions and approval rates were set by the

Date: 5th October-2026

author; the results therefore illustrate the mechanism and its approximate magnitude and are not an empirical validation. In particular, the very low FAR partly reflects the assumption that risk signals separate attackers from legitimate users reasonably well. (2) Adaptive adversaries who learn to keep their risk score low, or who compromise a trusted device, were not modelled. (3) Availability costs in the field, such as degraded connectivity that makes step-up or approval impossible, are only crudely represented by the 3% exceptional-use group. (4) The model does not address integrity (the Biba concern [2]), covert channels or cryptographic protection of stored documents. Future work should replay anonymised access logs from a real system, red-team the risk engine, and test usability with operators.

5. Conclusion

A hybrid access-control model that keeps a Bell–LaPadula lattice as a non-bypassable floor and adds ABAC need-to-know, device and network context, risk-adaptive thresholds and a two-person rule for the top classification removed nearly all simulated attack paths that defeat label-based and role-based control, at an availability cost of under 3% of legitimate requests. The result is consistent with current zero-trust direction, but it rests on a synthetic evaluation and should be treated as a design hypothesis to be validated on operational data before adoption.

Declarations

Data and code availability. The simulation was implemented in Python (NumPy) with ten fixed random seeds, so the reported tables are reproducible from the stated parameters. **Conflicts of interest.** None declared. **Funding.** None declared. [Edit these statements as required by the target journal.]

REFERENCES:

- [1] Bell DE, LaPadula LJ. Secure computer systems: mathematical foundations. Technical Report MTR-2547, Vol. I. Bedford (MA): The MITRE Corporation; 1973.
- [2] Biba KJ. Integrity considerations for secure computer systems. Technical Report MTR-3153. Bedford (MA): The MITRE Corporation; 1977.
- [3] Denning DE. A lattice model of secure information flow. Communications of the ACM. 1976;19(5):236–243.
- [4] Clark DD, Wilson DR. A comparison of commercial and military computer security policies. In: Proceedings of the IEEE Symposium on Security and Privacy; 1987. p. 184–194.
- [5] Brewer DFC, Nash MJ. The Chinese Wall security policy. In: Proceedings of the IEEE Symposium on Security and Privacy; 1989. p. 206–214.
- [6] Sandhu RS, Coyne EJ, Feinstein HL, Youman CE. Role-based access control models. IEEE Computer. 1996;29(2):38–47.
- [7] Ferraiolo DF, Sandhu R, Gavrila S, Kuhn DR, Chandramouli R. Proposed NIST standard for role-based access control. ACM Transactions on Information and System Security. 2001;4(3):224–274.

Date: 5th October-2026

- [8] Hu VC, Ferraiolo D, Kuhn R, Schnitzer A, Sandlin K, Miller R, Scarfone K. Guide to Attribute Based Access Control (ABAC) definition and considerations. NIST Special Publication 800-162. Gaithersburg (MD): National Institute of Standards and Technology; 2014. doi:10.6028/NIST.SP.800-162.
- [9] Rose S, Borchert O, Mitchell S, Connelly S. Zero Trust Architecture. NIST Special Publication 800-207. Gaithersburg (MD): National Institute of Standards and Technology; 2020. doi:10.6028/NIST.SP.800-207.
- [10] Servos D, Osborn SL. Current research and open problems in attribute-based access control. *ACM Computing Surveys*. 2017;49(4):Article 65.
- [11] Jin X, Krishnan R, Sandhu R. A unified attribute-based access control model covering DAC, MAC and RBAC. In: *Data and Applications Security and Privacy XXVI (DBSec 2012)*. Lecture Notes in Computer Science, vol. 7371. Berlin: Springer; 2012. p. 41–55.
- [12] Park J, Sandhu R. The UCON_ABC usage control model. *ACM Transactions on Information and System Security*. 2004;7(1):128–174.
- [13] Saltzer JH, Schroeder MD. The protection of information in computer systems. *Proceedings of the IEEE*. 1975;63(9):1278–1308.
- [14] Harrison MA, Ruzzo WL, Ullman JD. Protection in operating systems. *Communications of the ACM*. 1976;19(8):461–471.
- [15] Lampson BW. Protection. *ACM SIGOPS Operating Systems Review*. 1974;8(1):18–24.
- [16] Joint Task Force. Security and privacy controls for information systems and organizations. NIST Special Publication 800-53, Revision 5. Gaithersburg (MD): National Institute of Standards and Technology; 2020.
- [17] U.S. Department of Defense. Zero Trust Strategy. Washington (DC): Department of Defense; November 2022.
- [18] Bertino E, Sandhu R. Database security—concepts, approaches, and challenges. *IEEE Transactions on Dependable and Secure Computing*. 2005;2(1):2–19.
- [19] Anderson RJ. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 3rd ed. Indianapolis: Wiley; 2020.
- [20] Kuhn DR, Coyne EJ, Weil TR. Adding attributes to role-based access control. *IEEE Computer*. 2010;43(6):79–81.

