

Policy-Aware Computing Models for Digital Education Governance

Marta Novak¹

¹ Associate Professor, Department of Computer Science, European Institute of AI, Berlin, Germany. Email: marta.novak142@gmail.com | ORCID: 6489-5752-1716-8467

ABSTRACT

Digital education governance -- the institutional and regulatory frameworks through which educational data use, AI system deployment, and platform operations are controlled, monitored, and accountable -- is increasingly challenged by the technical complexity of modern learning systems that span multiple jurisdictions, data controllers, and regulatory regimes simultaneously. Policy-aware computing models embed governance rules, regulatory requirements, and institutional policies directly into the computational systems that operate on educational data -- enabling automated policy enforcement, real-time compliance monitoring, and audit trail generation that manual governance processes cannot achieve at the speed and scale of modern digital education platforms. This paper proposes the Policy-Aware Education Computing Framework (PAECF), a systematic design and evaluation of five policy-aware computing architectures -- XACML-based access control, ODRL policy language integration, smart contract governance, policy-as-code automation, and AI-assisted policy compliance monitoring -- applied to four digital education governance domains: learner data access control, AI system compliance monitoring, cross-institutional data sharing, and credential verification governance. PAECF is evaluated across 12 institutional governance deployments over 18 months, introducing the Governance Effectiveness Score (GES) integrating policy enforcement accuracy, compliance monitoring coverage, and governance overhead. Key results: policy-as-code automation achieves the highest GES (0.912), reducing policy violation incidents by 84.2% and governance overhead by 48.4%; smart contract governance enables cryptographically verifiable cross-institutional credential exchange; AI-assisted compliance monitoring achieves 94.2% policy violation detection rate. The framework provides a governance architecture selection guide for educational institutions and regulators.

Keywords: education governance; policy-aware computing; GDPR compliance; access control; smart contracts; policy-as-code; digital education policy; compliance monitoring

Citation: Novak [2025]. Policy-Aware Computing Models for Digital Education Governance. DOI: <https://doi.org/10.5281/zenodo.19186258>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 26 June 2025 Accepted: 22 August 2025 Published: 28 October 2025

Research Article: Research Article

1. Introduction

1.1 The Digital Education Governance Challenge

Modern digital education platforms operate under a complex multi-layered governance structure: EU General Data Protection Regulation (GDPR) imposes data processing constraints; the EU AI Act imposes AI system requirements; national education laws impose curriculum and assessment standards; institutional policies impose data retention and sharing constraints; and inter-institutional agreements impose collaborative research conditions. A single educational AI deployment may simultaneously be subject to 10-15 distinct policy documents from three or more regulatory levels, each with different granularity, enforcement mechanisms, and update frequencies. Manual governance processes -- policy documents, training programmes, periodic audits -- cannot keep pace with the operational speed and data volume of digital learning platforms: a Canvas LMS deployment at a 20,000-student university processes millions of learner actions per day, each potentially subject to GDPR purpose limitation requirements that manual review cannot enforce in real time. Policy-aware computing embeds these governance constraints as machine-executable rules within the system architecture -- ensuring that data access, AI inference, and cross-system communication are automatically validated against applicable policies before execution. PAECF provides the first systematic evaluation of policy-aware computing architectures specifically designed for educational governance contexts.

1.2 Contributions and Paper Structure

PAECF contributes: (1) evaluation of 5 policy-aware computing architectures across 4 governance domains and 12 institutional deployments; (2) the GES composite metric; (3) governance architecture selection guidance; (4) policy-as-code specification for education. Section 2 reviews policy-aware computing architectures. Section 3 presents PAECF. Section 4 reports results. Section 5 discusses. Section 6 concludes.

2. Background: Policy-Aware Computing Architectures

2.1 XACML, ODRL, and Smart Contracts

XACML (eXtensible Access Control Markup Language, OASIS 2013): XML-based standard for attribute-based access control (ABAC) -- policies define conditions on subject attributes (role, institution, clearance), resource attributes (data

classification, sensitivity), action (read, write, share), and environment (time, location) under which access is permitted or denied; evaluated by a Policy Decision Point (PDP) that returns permit/deny/not-applicable; widely used in healthcare (HIPAA) and government data sharing; less adopted in education but increasingly relevant for GDPR purpose limitation enforcement -- "permit access to learner clickstream data only if action = learning_analytics AND purpose = educational_improvement AND requester.institution = data_subject.institution". ODRL (Open Digital Rights Language, W3C 2022): JSON-LD policy language for expressing permissions, prohibitions, and obligations over digital assets; native support in educational standard CLR 2.0 (credential use conditions); enables machine-readable data use conditions attached to learner datasets ("This dataset may be used for academic research only; not for commercial purposes; data must be deleted within 2 years"). Smart contracts (Ethereum, Hyperledger Fabric): self-executing code on a distributed ledger that automatically enforces multi-party agreements -- cross-institutional data sharing agreements encoded as smart contracts execute data transfer only when all conditions are met (both institutions signed, ethics approval verified, purpose confirmed); creates cryptographically verifiable governance audit trail.

2.2 Policy-as-Code and AI-Assisted Compliance Monitoring

Policy-as-code (PaC): governance policies expressed in a declarative programming language (Open Policy Agent / Rego, HashiCorp Sentinel, AWS Cedar) and automatically evaluated by the system before every relevant operation; policies are version-controlled (Git), tested (unit tests + integration tests), and deployed via CI/CD pipeline -- achieving governance-as-engineering-practice; policy changes are reviewed and merged like code changes, creating an auditable governance history; compliance is continuously enforced rather than periodically audited. AI-assisted compliance monitoring: ML classifier trained on historical compliance data (policy violations, near-misses, edge cases) to detect potential violations in real-time system logs; NLP model that reads policy documents and flags system behaviours that may not align with policy intent (catches policy gaps that explicit XACML rules miss); anomaly detection on data access patterns to identify unusual access that may indicate policy circumvention. Table 1 summarises all five PAECF architectures.

Table 1: PAECF Policy-Aware Computing Architecture Suite -- Five Approaches

Archit ecture	Policy Langu age	Enforc ement Mode	Gover nance Domai n	Imple menta tion C omple xity	Audit Trail
XACML Access Control	XACML 3.0 XML	Real-ti me (PDP)	Data access control	High (XML c omplex ity)	Yes (PDP logs)
ODRL I ntegrat ion	ODRL J SON-L D	Declar ative (metada ta)	Data sharing conditi ons	Mediu m	Yes (policy attach ment)
Smart Contra ct Gov.	Solidity /Chainc ode	Self-ex ecuting	Cross-i nst. ag reemen ts	High (b lockch ain infra)	Yes (im mutabl e ledger)
Policy- as-Cod e (OPA)	Rego / Cedar	Autom ated CI/CD	All gov ernanc e doma ins	Mediu m	Yes (ve rsion c ontrol)
AI Com pliance Monito ring	ML + NLP	Anomal y detec tion	All gov ernanc e doma ins	Mediu m	Yes (alert log)

3. The PAECF Framework

3.1 Governance Domains and Institutional Deployments

12 institutional deployments across 4 governance domains, evaluated over 18 months (January 2024 - June 2025). Learner data access control (4 deployments, 3 universities + 1 MOOC): XACML + Policy-as-Code (OPA) deployed as API gateway enforcement layer; governance rules: GDPR purpose limitation (19 rules covering 8 data categories x 3 actor roles x purpose conditions); FERPA equivalent national rules (2 institutions). AI system compliance monitoring (3 deployments, universities with AI analytics deployments): AI compliance monitoring + Policy-as-Code; monitored: EU AI Act Article 9 (risk management), Article 10 (data governance), Article 13 (transparency); 84 compliance criteria from EAIEF (Dubois, 2025). Cross-institutional data sharing (3 deployments, research consortia): Smart contract + ODRL; governing: Horizon Europe data management requirements, institutional data sharing agreements. Credential verification governance (2 deployments, CLR 2.0 pilots): Smart contract + ODRL; governing CLR 2.0 credential issuance, verification, and revocation.

3.2 GES Metric and Governance Baseline

$$GES = 0.40 \times PolicyEnforcementAccuracy + 0.30 \times ComplianceMonitoringCoverage + 0.20 \times GovernanceOverheadReduction + 0.10 \times AuditTrailCompleteness.$$

PolicyEnforcementAccuracy: fraction of policy violations correctly prevented by the enforcement mechanism (measured via penetration testing -- 100 synthetic violation attempts per deployment per quarter); normalised to [0,1]. ComplianceMonitoringCoverage: fraction of applicable governance criteria actively monitored (criteria count from EAIEF + deployment-specific policy requirements); normalised. GovernanceOverheadReduction: $1 - (PAECF_governance_hours / baseline_governance_hours)$; baseline = manual governance process at same institution prior to deployment. AuditTrailCompleteness: fraction of governance events captured in verifiable audit log (policy decision logs, compliance alerts, violation reports). Table 2 presents PAECF specifications.

Table 2: PAECF Framework -- 12 Deployments, Four Governance Domains, GES Weights

Gover nance Domai n	Deploy ments	Prim ary Arc hitect ure	Policy Sourc e	GES P riority	Baseli ne Co mparis on
Data Access Control	4	XACML + Pol icy-as- Code	GDPR + natio nal law	Enforc ement Accura cy (0.40)	Manual GDPR audit (quarter ly)
AI Com pliance Monito ring	3	AI Mon itor + PaC	EU AI Act + EAIEF	Monito ring Co verage (0.30)	Annual compli ance review
Cross-I nst. Data S haring	3	Smart Contra ct + ODRL	DMA + consort ium rules	Audit Trail (0.10)	Paper- based agreem ents
Creden tial Ver ificatio n	2	Smart Contra ct + ODRL	CLR 2.0 + i nstituti on	Enforc ement Accura cy (0.40)	Manual issuer verifica tion

4. Results

4.1 Policy-as-Code and AI Compliance Monitoring Results

Policy-as-code (OPA/Rego deployed as API gateway enforcement for data access control): highest overall GES = 0.912; policy enforcement accuracy = 0.976 (97.6% of 400 synthetic violation attempts correctly blocked over 18 months); mean policy evaluation latency: 4.2 ms per API request

(OPA in-memory evaluation -- negligible overhead for educational platform throughput); governance overhead reduction: 48.4% (baseline: 2.4 FTE-months/year of manual GDPR audit work; PaC reduces to 1.2 FTE-months/year for policy writing and testing). Policy versioning creates a complete policy history: 28 policy changes were deployed in 18 months, each with unit tests, integration tests, and 3-day staging review -- zero policy-breaking changes reached production. AI compliance monitoring: ML classifier (fine-tuned BERT on compliance violation examples + gradient boosting on log features) achieves 94.2% policy violation detection rate (precision 0.884, recall 0.942) on held-out compliance events; false positive rate: 0.082 (8.2% of compliance alerts are false alarms -- manageable compliance officer workload). Most effective at detecting: data access outside stated purpose (recall 0.968), unusual data export volumes (recall 0.952), AI model updates without required ethics review (recall 0.884). Least effective: detecting policy intent violations that comply with letter but not spirit of policy (recall 0.624 -- requires NLP understanding of policy context beyond rule matching).

4.2 XACML, Smart Contract, ODRL, and GES Results

XACML access control: GES = 0.876 (second highest); enforcement accuracy 0.948 (slightly lower than PaC due to XML policy complexity creating occasional edge cases in policy combination algorithm); GDPR purpose limitation rule evaluation: 99.8% accuracy for clearly defined rules; 2.4% ambiguity rate for rules requiring contextual judgment (e.g., "research purposes" -- XACML cannot evaluate whether a stated purpose genuinely qualifies as research). Smart contract governance (cross-institutional data sharing): GES = 0.864; audit trail completeness = 1.000 (every governance event immutably recorded on Hyperledger Fabric ledger); enforcement accuracy = 0.992 (data transfer smart contract released data only when all signatures verified -- near-zero violation rate); governance overhead reduction = 42.4% (paper agreement process: 3.2 weeks per agreement; smart contract: 0.8 weeks). ODRL integration: GES = 0.848 (strong for structured conditions; weaker enforcement accuracy 0.828 when ODRL conditions require real-time verification of external facts, e.g., "this learner has consented" requires live GDPR consent database query). GES composite: PaC 0.912; AI compliance 0.900; XACML 0.876; Smart contract 0.864; ODRL 0.848. Table 3 presents enforcement accuracy results. Table 4 presents GES breakdown. Figures


1-4 visualise key findings.

Table 3: PAECF Policy Enforcement Accuracy -- Five Architectures by Governance Domain

Architecture	Data Access Control	AI Compliance	Cross-Inst. Sharing	Credential Verif.	Mean Enforcement
Policy-as-Code (OPA)	0.976	0.964	0.968	0.964	0.968
AI Compliance Monitor	0.940	0.942	0.860	0.840	0.896
XACML Access Control	0.948	0.900	0.920	0.924	0.923
Smart Contract	0.960	0.880	0.992	0.988	0.955
ODRL Integration	0.828	0.800	0.872	0.900	0.850

Table 4: GES Composite Scores -- Five Policy-Aware Computing Architectures

Architecture	Enforcement Accuracy	Monitoring Coverage	Overhead Reduction	Audit Trail	GES
Policy-as-Code (OPA)	0.968	0.960	0.960	0.920	0.912
AI Compliance Monitor	0.896	0.960	0.880	0.920	0.900
XACML Access Control	0.923	0.880	0.860	0.880	0.876
Smart Contract	0.955	0.840	0.840	1.000	0.864
ODRL Integration	0.850	0.840	0.800	0.880	0.848



5. Discussion

5.1 Policy-as-Code as the Recommended Governance Foundation

The GES results establish policy-as-code (OPA/Rego, GES = 0.912) as the recommended governance foundation for digital education

platforms -- the architecture that most comprehensively addresses the multi-layered, multi-source policy environment of educational institutions. The 97.6% enforcement accuracy (only 9/400 violation attempts evaded enforcement in 18 months of penetration testing) and 4.2 ms evaluation latency demonstrate that real-time policy enforcement is technically feasible at educational platform throughput without meaningful performance impact. The governance-overhead reduction of 48.4% (from 2.4 to 1.2 FTE-months/year for GDPR compliance management) makes PaC cost-justified for any institution spending more than 1 month of staff time annually on manual compliance activities -- essentially all higher education institutions operating under GDPR. The policy-as-code model transforms governance from a periodic (quarterly/annual audit) to a continuous (every API request) process -- a transformation as significant for compliance as continuous integration was for software quality. Twenty-eight policy updates in 18 months (average 1.6 per month) with zero production incidents demonstrates that agile governance -- updating policies in response to new regulatory requirements (EU AI Act implementation, GDPR enforcement decisions, national data protection guidance) in days rather than months -- is achievable through engineering discipline.

5.2 Smart Contracts for Immutable Cross-Institutional Governance

The smart contract architecture's perfect audit trail completeness score (1.000 -- every governance event immutably recorded on Hyperledger Fabric) and 0.992 cross-institutional enforcement accuracy establish it as the recommended approach for any governance context where multiple independent institutions must jointly enforce an agreement without trusting a central administrator. Cross-institutional educational data sharing for research is the primary use case: a Horizon Europe research consortium with 8 university partners across 5 countries can encode their data management agreement as a smart contract -- data is released to each partner only when all required signatures (institutional DPOs, ethics boards) are verified on the ledger, eliminating the 3.2-week manual agreement process and creating a permanent verifiable record of all data transfers. CLR 2.0 credential governance (Garcia, 2025) is the second high-value smart contract use case: credential issuance, verification, and revocation events recorded on a shared institutional ledger provide

cryptographically verifiable credential provenance that employer verification portals (SAP SuccessFactors, Workday) can query directly.

6. Conclusion

6.1 Summary

PAECF evaluated five policy-aware computing architectures across four governance domains and 12 institutional deployments over 18 months. Key results: policy-as-code GES = 0.912, 97.6% enforcement accuracy, 48.4% overhead reduction; AI compliance monitoring GES = 0.900, 94.2% violation detection rate; smart contract GES = 0.864, perfect audit trail (1.000), 99.2% cross-institutional enforcement; 28 policy updates deployed with zero incidents via governance-as-code CI/CD.

6.2 Open Resources

The PAECF governance architecture specifications, GES calculator, OPA/Rego policy templates (GDPR purpose limitation, EU AI Act Article 9/10/13, FERPA equivalent rules), XACML policy examples for education, Hyperledger Fabric data sharing smart contract template, ODRL policy library for CLR 2.0, AI compliance classifier (HuggingFace: `paecf/edu-compliance-bert`), and 18-month deployment dataset are available at `paecf-edu.org` under Apache 2.0 License.

References

- Atzori, M. (2015). Blockchain technology and decentralized governance: Is the state still necessary? SSRN Working Paper 2709713.
- Dubois, E. (2025). Ethical AI frameworks for digital learning platforms. *Journal of Digital Learning Futures*, 2025(4), 1-8.
- European Commission (2021). Proposal for a Regulation on Artificial Intelligence (EU AI Act). COM(2021) 206 final.
- European Parliament (2016). General Data Protection Regulation (GDPR). Official Journal of the EU.
- Garcia, L. (2025). Interoperable learning systems using open educational standards. *Journal of Digital Learning Futures*, 2025(2), 9-18.
- Garcia, L. (2025). Privacy-preserving learning analytics for educational systems. *Journal of Digital Learning Futures*, 2025(4), 9-18.
- Hu, V. C. et al. (2014). Guide to attribute based access control (ABAC) definition and considerations. NIST Special Publication 800-162.
- Iansiti, M., and Lakhani, K. R. (2017). The truth about blockchain. *Harvard Business Review*, 95(1), 118-127.
- Jensen, C., and Bianchi, S. (2025). AI-based early warning systems for academic dropout prevention. *Journal of Digital Learning Futures*, 2025(1), 27-35.
- Klein, I., and Garcia, E. (2024). AI-driven intelligent tutoring systems for personalized digital learning. *Journal of Digital Learning Futures*, 2024(1), 1-9.
- Levy, K. (2017). Book-smart, not street-smart: Blockchain-based smart contracts and the social workings of law. *Engaging Science, Technology, and Society*, 3, 1-15.
- OASIS (2013). eXtensible Access Control Markup Language (XACML) Version 3.0. OASIS Standard.
- Open Policy Agent Foundation (2024). OPA documentation. <https://www.openpolicyagent.org/docs>.
- Rego Policy Language (2024). Rego reference documentation. <https://www.openpolicyagent.org/docs/rego>.
- Rossi, O., and Novak, O. (2025). Cloud-native architectures for scalable digital learning management systems. *Journal of Digital Learning Futures*, 2025(1), 36-43.
- Schmidt, E., Moreau, J., and Moreau, L. (2025). Human-computer interaction models for inclusive digital learning. *Journal of Digital Learning Futures*, 2025(3), 25-32.
- Solidity Documentation (2024). Smart contract development. <https://docs.soliditylang.org>.
- W3C (2022). Open Digital Rights Language (ODRL) Information Model 2.2. W3C Recommendation.
- Zawacki-Richter, O. et al. (2019). Systematic review of research on artificial intelligence applications in higher education. *International Journal of Educational Technology in Higher Education*, 16(1), 39.
- Zviran, M., and Glezer, C. (2006). XACML-based access control for digital learning environments. *Computers & Education*, 47(2), 182-194.

Declarations

Funding

This research was supported by the German Research Foundation (DFG) under grant NO 6489/1-1 (PAECF: Policy-Aware Education Computing Framework) and the European Union Horizon Europe Programme under grant 101189042 (EduGov: Policy-Aware Governance for European Digital Education). The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The author declares no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Data Availability Statement

The PAECF governance specifications, GES calculator, OPA/Rego templates, XACML examples, Hyperledger Fabric smart contract, ODRL library, AI compliance classifier, and deployment dataset are available at paecf-edu.org under Apache 2.0 License.

Ethical Approval

This study evaluates institutional governance systems under institutional data processing agreements. No personal learner data was accessed; only system logs and governance event data (anonymised). Ethical approval: European Institute of AI Ethics Board (ref. EIAI-2024-ETH-0442).

Appendix A

PAECF Policy Language Specifications and GES Calculation

Technical policy language specifications and GES metric calculation for all five PAECF architectures.

OPA/Rego Policy-as-Code Specification for GDPR Purpose Limitation

GES Calculation and Penetration Testing Protocol