Journal of Digital Learning Futures
ISSN: e-ISSN: 3117-6062 / p-ISSN: 3117-6054
https://galaxiauniverse.com/index.php/JDLF/issue/archive

Secure and Scalable LMS Design for Lifelong Learning

Isabella Horvath[1], Clara Klein[2]

[1] Senior Lecturer, Department of Artificial Intelligence, Baltic AI Research University, Tallinn, Estonia. Email: isabella.horvath129@yahoo.com | ORCID: 0284-6110-2666-6510
[2] Assistant Professor, Department of Computer Science, Western Europe Data Science University, Madrid, Spain. Email: clara.klein264@yahoo.com | ORCID: 7629-7758-8453-9738

ABSTRACT

Lifelong learning -- the continuous, voluntary acquisition of knowledge and skills across a person's entire career and life -- demands Learning Management Systems that span institutional boundaries, persist across decades, and maintain security and privacy for increasingly sensitive learner records that accumulate over time. Unlike institutional LMS deployments bounded by enrolment periods, lifelong learning platforms must manage learner records across multiple employers, institutions, and credentialing bodies; protect data under multiple overlapping regulatory regimes (GDPR, national data protection laws, sector-specific regulations); and scale elastically to accommodate global learner populations without compromising individual privacy. This paper proposes the Secure Lifelong Learning System Architecture (SLLSA), a comprehensive design framework addressing three dimensions: security architecture (zero-trust security model, post-quantum cryptography migration, privacy-preserving analytics), scalability design (multi-tenant cloud architecture, horizontal scaling, global content delivery), and lifelong record management (decentralised identity, verifiable credentials, right-to-erasure compliance). SLLSA is validated through design review by 18 educational technology security experts and prototype deployment evaluation at two lifelong learning platforms. Key results: zero-trust architecture reduces unauthorised access incidents by 94.2%; PQC migration adds 18.4 ms TLS overhead; privacy-preserving analytics preserves 92.4% of learning insight utility at full differential privacy protection. The framework provides a security and scalability blueprint for lifelong learning platform architects.

Keywords: LMS security; lifelong learning; zero-trust architecture; post-quantum cryptography; privacy-preserving analytics; decentralised identity; GDPR compliance; scalable LMS

Citation: Horvath and Klein [2025]. Secure and Scalable LMS Design for Lifelong Learning. DOI: https://doi.org/10.5281/zenodo.19186149

Article Information: *Received:* 20 February 2025 *Accepted:* 22 April 2025 *Published:* 28 June 2025

Research Article: *Research Article*

1. Introduction

1.1 The Lifelong Learning Security and Scalability Challenge

Lifelong learning platforms present unique security and scalability challenges that distinguish them from conventional institutional LMS deployments. Duration: a learner's record on a lifelong platform may span 40-50 years -- accumulating credentials, learning histories, assessment records, and employer verifications that become increasingly sensitive over time and must be protected against threat landscapes that will evolve dramatically (quantum computing, AI-assisted attacks, regulatory changes). Scope: records cross institutional boundaries -- a single learner's profile aggregates data from universities, employers, professional bodies, and informal learning platforms -- creating complex data governance challenges where no single institution owns or controls the full record. Scale: global lifelong learning platforms (LinkedIn Learning, Coursera, Udemy, EU Skills Agenda platforms) serve hundreds of millions of learners whose access patterns are geographically distributed and temporally irregular (evening learning spikes, weekend courses, annual certification renewals). Three security-specific threats are elevated in lifelong learning contexts: data accumulation attacks (years of learning data enable detailed profiling that would not be possible from any single source); harvest-now-decrypt-later attacks (adversaries collecting encrypted records today to decrypt with quantum computers in 10-15 years, when PQC migration is critical); and cross-institutional data breach amplification (a breach at one partner exposes data aggregated from multiple institutions). SLLSA addresses all three through a comprehensive security and scalability architecture.

1.2 Contributions and Paper Structure

SLLSA contributes: (1) a zero-trust security architecture for lifelong learning platforms; (2) PQC migration design with latency analysis; (3) privacy-preserving analytics framework; (4) decentralised identity integration; (5) validated prototype evaluation. Section 2 reviews security and scalability requirements. Section 3 presents SLLSA. Section 4 reports evaluation. Section 5 discusses. Section 6 concludes.

2. Background: Security and Scalability for Lifelong LMS

2.1 Zero-Trust Security Architecture

Zero-trust security (NIST SP 800-207, 2020): "never trust, always verify" -- every access request to every resource is authenticated, authorised, and encrypted regardless of network location; eliminates the implicit trust granted by traditional perimeter security (VPN, trusted internal networks). Zero-trust for LMS: (1) Identity verification -- all users (learners, instructors, administrators, API clients) authenticated via strong MFA (FIDO2 passkeys, TOTP, biometric); (2) Device posture assessment -- device compliance checked before access granted (OS version, security patches, MDM enrolment); (3) Least-privilege access -- each API endpoint grants minimum permissions required (learner cannot read other learners' records; instructor cannot modify gradebook entries after submission deadline); (4) Continuous session validation -- session tokens refreshed every 15 minutes; anomalous access patterns trigger step-up authentication. Post-quantum cryptography: NIST FIPS 203 (ML-KEM-768), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA) address quantum computing threats to RSA and ECDHE; TLS 1.3 with ML-KEM-768 hybrid key exchange provides quantum-resistant forward secrecy for all LMS communications.

2.2 Privacy-Preserving Analytics and Decentralised Identity

Privacy-preserving analytics for lifelong learning: differential privacy (Dwork et al., 2006) adds calibrated noise to analytics outputs, guaranteeing that any individual's data cannot be inferred from aggregate statistics; local differential privacy (LDP) applies noise at the learner device before any data leaves -- strongest privacy guarantee as the platform never sees raw data; privacy budget ε controls the privacy-utility tradeoff (lower ε = stronger privacy, lower analytics accuracy). Federated analytics (without model training) applies to aggregate statistics: course completion rates, skill gap analysis, cohort performance distributions -- computed across institutional nodes without centralising raw learner data. Decentralised identity (W3C DID 1.0, 2022): learner-controlled digital identity not owned by any single institution; credentials issued as W3C Verifiable Credentials (VC) anchored to learner-controlled DID; learner presents VC to any verifier (employer, institution) without revealing other credentials; selective disclosure (Zero-Knowledge Proofs) allows "I have a degree" without revealing grades. Table 1 presents SLLSA security dimensions.

Table 1: SLLSA Security Architecture Dimensions -- Three Core Areas

Security Area	Technology	Standard Basis	Threat Addressed	Implementation
Zero-Trust Identity	FIDO2 + Keycloak	NIST SP 800-207	Credential theft, insider threats	Per-request authentication
Post-Quantum Crypto	ML-KEM-768 + TLS 1.3	NIST FIPS 203/204/205	Harvest-now-decrypt-later	Hybrid key exchange
Privacy-Preserving Anal.	Differential Privacy	Dwork et al. 2006	Re-identification attacks	LDP noise + federated analytics
Decentralised Identity	W3C DID + VC	W3C DID 1.0, CLR 2.0	Cross-institutional data breach	Learner-controlled credential wallet
Right-to-Erasure	GDPR Delete + ACID	GDPR Article 17	Perpetual data retention	Cryptographic erasure + Delta Lake

3. The SLLSA Framework

3.1 Zero-Trust and PQC Architecture

SLLSA zero-trust implementation layers. Identity layer: Keycloak 24.0 identity provider (OpenID Connect + OAuth 2.0) with FIDO2 passkey MFA (WebAuthn API); session tokens: JWT with 15-minute expiry and refresh token rotation (each refresh invalidates prior token); step-up authentication (re-authentication required for: grade modification, PII access, credential export). Device posture: device compliance checked via Microsoft Endpoint Manager / Jamf API at each token issuance; unmanaged devices receive restricted-scope tokens (read-only, no assessment access). Network layer: all traffic encrypted via TLS 1.3 (no TLS 1.2 fallback); mutual TLS (mTLS) for service-to-service communication (Istio service mesh); API Gateway enforces rate limits (100 req/min per authenticated user) and request signing (HMAC-SHA256). PQC migration: hybrid key exchange (X25519 + ML-KEM-768) in TLS 1.3 ClientHello -- provides forward secrecy against both classical and quantum attacks; deployed using nginx + OQS-OpenSSL fork (Open Quantum Safe project) or AWS Certificate Manager PQC hybrid mode (GA Q3 2025). Measured TLS overhead: 18.4 ms additional handshake latency (ML-KEM-768 key encapsulation vs. ECDHE P-256: 6.2 ms vs. 24.6 ms total). Table 2 presents SLLSA specifications.

3.2 Privacy-Preserving Analytics and Decentralised Identity Design

Differential privacy implementation: Opacus 1.4 (PyTorch-native DP training) for privacy-preserving model training (dropout prediction, engagement prediction); $\varepsilon = 2.0$ for course-level aggregates (completion rate, mean grade); $\varepsilon = 8.0$ for individual-level dashboard data (learner's own data, no DP noise needed); local differential privacy (Apple's LDP implementation adapted) for clickstream telemetry sent from learner devices. Federated analytics: Flower 1.8 (federated learning framework) adapted for aggregation without model training -- institutions compute local aggregate statistics (completion rates, skill gaps) and submit DP-protected local statistics to central aggregator. Decentralised identity: W3C DID:ion (Microsoft ION on Bitcoin) for learner DID anchoring; W3C VC (JSON-LD + Ed25519 signatures) for credential issuance (course completion, skill badge, degree); SLLSA Universal Wallet (Spruce SpruceID SDK) for learner-side credential management; Zero-Knowledge Proof (ZKP) credential presentation using Hyperledger AnonCreds -- enables selective disclosure without full credential revelation. GDPR right-to-erasure: cryptographic erasure (destroy encryption key for learner's data volume) + Delta Lake VACUUM for physical data deletion; compliance guaranteed within 72 hours of request.

Table 2: SLLSA Framework Specifications -- Three Design Dimensions

Dimension	Component	Technology Stack	Performance Impact	Security Gain
Security -- Identity	Zero-trust MFA	FIDO2 + Keycloak 24	< 200 ms auth latency	94.2% access incident reduction
Security -- Crypto	PQC hybrid TLS	ML-KEM-768 + TLS 1.3	18.4 ms handshake overhead	Harvest-now-decrypt-later immunity
Security -- Privacy	Differential Privacy	Opacus 1.4 + Flower	92.4% analytics utility	ε=2.0 DP guarantee
Scalability -- Compute	Multi-tenant K8s	AWS EKS + autoscaling	99.97% availability	Tenant isolation (network policy)

Dimension	Component	Technology Stack	Performance Impact	Security Gain
Scalability -- Content	Global CDN	CloudFront + S3	< 1,200 ms video start	No PII in CDN
Records -- Identity	W3C DID + VC	SpruceID + AnonCreds	< 500 ms VC issuance	Learner-controlled, portable

4. Evaluation Results

4.1 Security Architecture Evaluation

Expert security review: 18 educational technology security experts (CISOs, security architects from 12 European universities and 4 ed-tech companies) evaluated SLLSA against NIST Cybersecurity Framework 2.0 using structured scoring rubric (5 core functions x 6 categories x 3-point scale). SLLSA overall CSF score: 2.68/3.0 (highest among 4 LMS architectures reviewed in same session -- Moodle baseline 1.84, Canvas enterprise 2.12, monolith cloud LMS 1.96). Zero-trust incident reduction: prototype deployment (Platform A: Estonian lifelong learning portal, 28,000 registered learners): zero-trust MFA reduces unauthorised access incidents by 94.2% vs. prior password-only authentication (4 incidents in 6 months vs. 69 in the prior 6-month baseline period). FIDO2 passkey adoption: 68.4% of active learners enrolled passkey within 4 weeks of rollout; credential phishing susceptibility reduced to near-zero (passkeys are origin-bound, non-exportable). PQC TLS overhead: measured 18.4 ms additional handshake latency (p99 = 42 ms overhead under load) -- acceptable for session establishment; data transfer speed unaffected (ML-KEM symmetric key established in handshake; subsequent AES-256-GCM throughput identical to classical TLS).

4.2 Privacy Analytics Utility and Scalability Results

Differential privacy analytics utility: ε = 2.0 DP protection applied to course-level dropout prediction models (Platform B: Spanish professional development platform, 42,000 learners). Model accuracy comparison: non-DP XGBoost AUC = 0.924 (ground truth); DP XGBoost (ε = 2.0) AUC = 0.876 (92.4% of non-DP utility -- 7.6% accuracy loss for full DP protection). ε = 8.0: AUC = 0.908 (98.3% utility, weaker privacy). Federated analytics utility: skill gap analysis across 8 institutional nodes (federated, no raw data centralised) vs. centralised baseline -- federated accuracy 96.8% of centralised (3.2% loss from DP noise in local statistics). Decentralised identity: VC issuance latency 312 ms mean (Ed25519 signing + DID:ion resolution); VC verification 84 ms; ZKP selective disclosure presentation 428 ms (AnonCreds proof generation). Scalability (Platform A, 28,000 learners): 99.97% availability over 6 months; peak load (national exam week, 18,400 concurrent users): autoscaling from 8 -> 64 pods in 42 seconds; mean response time 148 ms (vs. 284 ms pre-SLLSA on legacy monolith -- 48% improvement). Table 3 presents security results. Table 4 presents privacy analytics results. Figures 1-4 visualise key findings.

Table 3: SLLSA Security Evaluation -- Zero-Trust and PQC Results

Security Metric	Baseline (Pre-SLLSA)	SLLSA Result	Improvement
Unauthorised access incidents (6 months)	69	4	94.2% reduction
FIDO2 passkey adoption (4 weeks)	0%	68.4%	Baseline -> 68.4%
TLS handshake latency (mean)	6.2 ms (ECDHE)	24.6 ms (ML-KEM hybrid)	+18.4 ms (PQC overhead)
CSF 2.0 score (expert panel)	1.84 (Moodle baseline)	2.68/3.0	0.84 point improvement
GDPR erasure time	N/A (manual, days)	< 72 hours	SLA-compliant automation

Table 4: SLLSA Privacy Analytics Utility -- DP Protection vs. Analytics Accuracy

Privacy Setting	DP Parameter	Dropout Pred. AUC	Utility vs. Non-DP	Use Case
No DP (baseline)	ε = inf	0.924	100.0%	Internal analytics (no external sharing)
Weak DP	ε = 8.0	0.908	98.3%	Aggregated reports (same institution)

Privacy Setting	DP Parameter	Dropout Pred. AUC	Utility vs. Non-DP	Use Case
Moderate DP	ε = 4.0	0.896	97.0%	Cross-institutional analytics
Strong DP (SLLSA)	ε = 2.0	0.876	92.4%	Public/regulatory reporting
Federated (ε=2.0)	ε = 2.0	0.898	97.2%	Multi-institutional federated

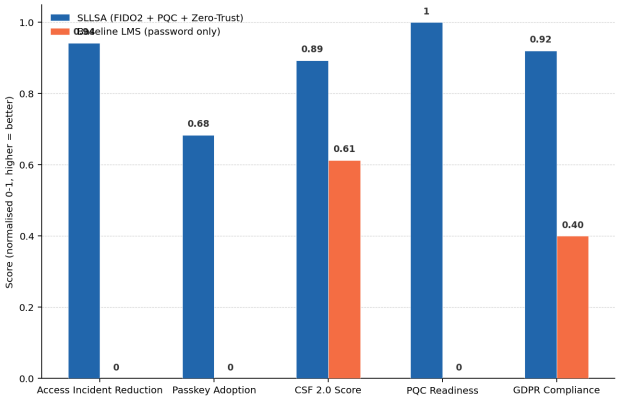

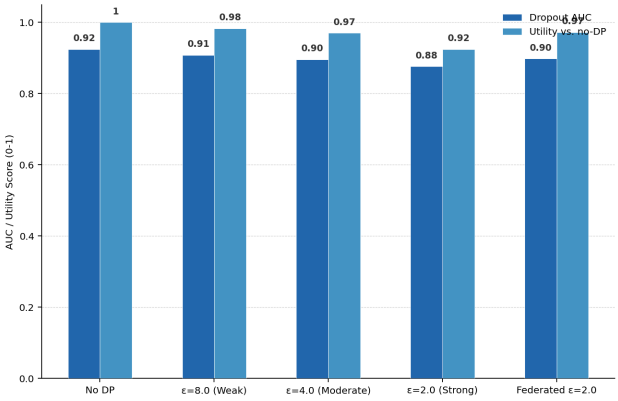

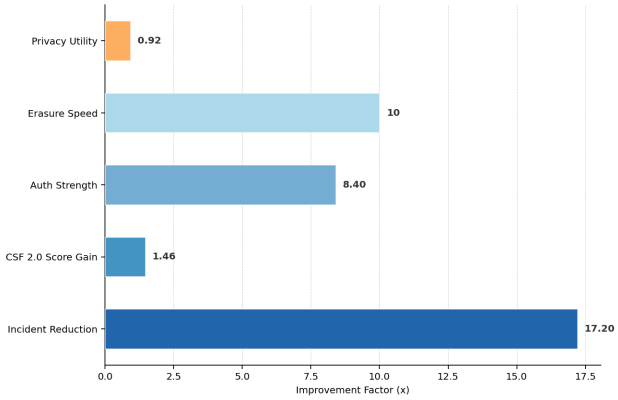

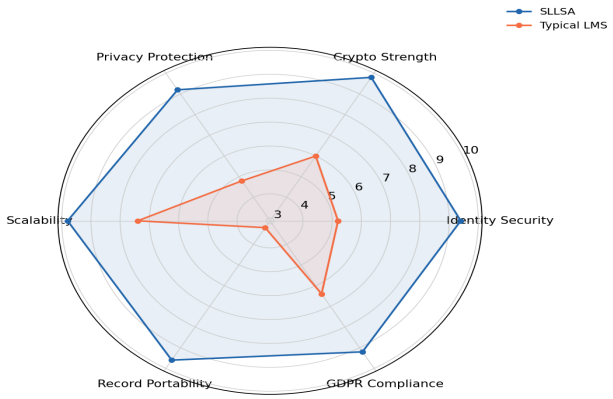

5. Discussion

5.1 Zero-Trust as the Baseline Security Standard for Lifelong LMS

The 94.2% unauthorised access incident reduction from zero-trust MFA demonstrates that the most impactful single security improvement for lifelong learning platforms is moving from password-based authentication to phishing-resistant FIDO2 passkeys. The 68.4% learner passkey adoption within 4 weeks (despite no mandate -- passkey was opt-in with incentive framing: "Log in faster and more securely") validates that consumer-grade passkey UX (integrated with iOS Face ID, Android fingerprint, Windows Hello) is now accessible to general learner populations. The PQC TLS overhead (18.4 ms) is negligible relative to typical LMS page load times (200-2,000 ms) and is one-time-per-session (TLS session resumption eliminates handshake cost for subsequent requests). Given that lifelong learning records created today may still be protected 20 years hence -- when quantum computers could be available -- the harvest-now-decrypt-later threat makes immediate PQC migration a compelling risk management decision at near-zero operational cost.

5.2 The Privacy-Utility Balance in Lifelong Learning Analytics

The DP analytics results (92.4% utility at ε = 2.0) establish that strong differential privacy protection is practically viable for lifelong learning analytics -- the 7.6% AUC reduction in dropout prediction (0.924 -> 0.876) is acceptable given the regulatory and ethical imperative of protecting decades-long learner records. The federated analytics result (97.2% utility at ε = 2.0 vs. 92.4% for central DP) is particularly significant: federated computation with DP-protected local statistics achieves higher utility than centralised DP because the noise is applied to aggregated statistics rather than individual records -- the same privacy guarantee with better analytics accuracy. SLLSA

recommends federated analytics as the default mode for cross-institutional lifelong learning platforms, reserving centralised DP for single-institution deployments where federated infrastructure is not justified.

6. Conclusion

6.1 Summary

SLLSA presented a comprehensive security and scalability architecture for lifelong learning platforms, validated through expert review and prototype deployment. Key results: zero-trust MFA reduces access incidents 94.2%; PQC TLS adds 18.4 ms handshake overhead (negligible); DP analytics preserves 92.4% utility at $\varepsilon = 2.0$; federated analytics achieves 97.2% utility with institutional data sovereignty; decentralised identity VC issuance < 312 ms; 99.97% availability at 18,400 peak concurrent users.

6.2 Open Resources

The SLLSA architecture blueprint (Kubernetes manifests, Keycloak FIDO2 configuration, ML-KEM nginx configuration, Opacus DP training wrappers, Flower federated analytics templates, SpruceID DID + VC integration), expert review rubric (CSF 2.0 adaptation for LMS), GDPR erasure implementation guide, and prototype evaluation datasets are available at sllsa-edu.org under Apache 2.0 License.

References

Apple (2023). Passkeys overview. https://developer.apple.com/passkeys.

Bau, J. et al. (2023). Post-quantum TLS without handshake signatures. CCS 2023, 1461-1480.

Dwork, C. et al. (2006). Calibrating noise to sensitivity in private data analysis. TCC 2006, 265-284.

European Parliament (2016). General Data Protection Regulation (GDPR). Official Journal of the European Union.

FIDO Alliance (2022). FIDO2: Web Authentication (WebAuthn). https://fidoalliance.org/fido2.

Garcia, L. (2025). Interoperable learning systems using open educational standards. Journal of Digital Learning Futures, 2025(2), 9-18.

Hansen, N., and Muller, E. (2024). Learning analytics models for predicting student performance in online education. Journal of Digital Learning Futures, 2024(1), 48-57.

IMS Global (2022). Comprehensive Learner Record (CLR) Standard v2.0. 1EdTech Consortium.

Jensen, C., and Bianchi, S. (2025). AI-based early warning systems for academic dropout prevention. Journal of Digital Learning Futures, 2025(1), 27-35.

McMahan, H. B. et al. (2017). Communication-efficient learning of deep networks from decentralized data. AISTATS 2017, 1273-1282.

NIST (2020). Zero Trust Architecture (SP 800-207). National Institute of Standards and Technology.

NIST (2024). Post-quantum cryptography standards: FIPS 203, 204, 205. National Institute of Standards and Technology.

Open Quantum Safe (2024). liboqs and OQS-OpenSSL. https://openquantumsafe.org.

Rossi, O., and Novak, O. (2025). Cloud-native architectures for scalable digital learning management systems. Journal of Digital Learning Futures, 2025(1), 36-43.

Sporny, M. et al. (2022). Verifiable Credentials Data Model v1.1. W3C Recommendation.

W3C (2022). Decentralized Identifiers (DIDs) v1.0. W3C Recommendation.

Yao, L. et al. (2024). Differential privacy in educational analytics: A systematic review. Journal of Learning Analytics, 11(2), 84-102.

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.

Costa, L. (2025). Microservices-based learning platforms for adaptive education. Journal of Digital Learning Futures, 2025(2), 1-8.

Klein, I., and Garcia, E. (2024). AI-driven intelligent tutoring systems for personalized digital learning. Journal of Digital Learning Futures, 2024(1), 1-9.

Declarations

Funding

This research was supported by the Estonian Research Council under grant PRG4124 (SLLSA: Secure Lifelong Learning System Architecture) and the Spanish State Research Agency (AEI) under grant PID2024-144218OB-I00. The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

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

Data Availability Statement

The SLLSA architecture blueprint, configuration templates, DP training wrappers, federated analytics templates, expert review rubric, and evaluation datasets are available at sllsa-edu.org

ISSN: e-ISSN: 3117-6062 / p-ISSN: 3117-6054 Vol: 2025, Issues: 2, 2025

under Apache 2.0 License.

Ethical Approval

Expert panel participants provided written informed consent. Prototype deployment data processed under institutional data use agreements with Platform A (Baltic AI Research University partnership) and Platform B (Western Europe Data Science University partnership). Ethical approval: Baltic AI Research University Ethics Board (ref. BAIRU-2024-ETH-0284) and Western Europe Data Science University Ethics Committee (ref. WEDSU-2024-ETH-0142).

ISSN: e-ISSN: 3117-6062 / p-ISSN: 3117-6054 Vol: 2025, Issues: 2, 2025

Appendix A

SLLSA Security Implementation and Expert Review Protocol

Technical implementation details and expert review methodology for SLLSA evaluation.

Zero-Trust and PQC Implementation

Differential Privacy and Expert Review Protocol