

Cloud Computing Security and Risk Management in Smart Learning Environments: An Enterprise Risk Management Framework for Institutional Adoption

Kanita Haider ¹, Danish Mahmud ², Oishe Al Mariz ³

¹ Chittagong University of Engineering and Technology, Chittagong, Bangladesh

² University of the Cumberland, Kentucky, USA

³ Chittagong University of Engineering and Technology, Chittagong, Bangladesh.

*Corresponding author: kanita.haider4422@gmail.com

Received 15 January 2026; Revised 23 April 2026; Accepted 19 May 2026; Published: 25 July 2026

KEYWORDS

Cloud computing security
Enterprise risk management
Smart learning environments
Analytic hierarchy process
Monte Carlo simulation
Vendor risk management
Higher education cybersecurity

Abstract

As institutions migrate learning management systems, e-learning platforms, and learning-analytics infrastructure to third-party cloud environments, they inherit cloud-specific cybersecurity and data-privacy risks that education-sector security research has documented extensively, phishing, insider threats, vendor risk, and gaps in breach disclosure, but has not yet paired with a formal enterprise-grade risk-management response. This paper argues that Enterprise Risk Management (ERM), and specifically the baseline cloud-security requirements developed for general enterprise deployments, offers a more mature and transferable governance model than education-sector practice has produced on its own. Synthesizing the higher-education threat literature with the enterprise ERM literature, the paper develops a threat-to-control relevance matrix, a four-tier data classification for smart learning environment (SLE) data, a three-lines governance layer, and a dedicated vendor-risk component. The framework is then evaluated with a hybrid quantitative design: an analytic hierarchy process (AHP) prioritizes seven components (consistency ratio 0.008); a FAIR-style Monte Carlo model with 100,000 trials estimates annual loss under inherent, baseline, and ERM scenarios; an independent-cascade model of a twelve-service dependency network quantifies systemic exposure; and stepwise return-on-investment and sensitivity analyses translate the results into an adoption sequence. Under illustrative, literature-informed parameters, the ERM scenario lowers expected annual loss from \$4.15M to \$1.36M (67% lower), reduces the probability of exceeding a \$5 million annual loss from 26.2% to 2.6%, and cuts the probability of a cascade affecting four or more services from 23.1% to 3.8%. Governance consistency and control strength dominate output uncertainty. The paper concludes that institutional cybersecurity for smart learning environments should be governed as a strategic, leadership-owned risk function rather than delegated to IT as a technical afterthought, and it proposes practical recommendations for pre-migration risk assessment, executive risk ownership, and mandatory vendor risk documentation ahead of EdTech procurement.

1. Introduction

Higher education institutions (HEIs) have rapidly shifted the infrastructure that underpins teaching and learning, learning management systems (LMS), e-learning delivery platforms, and learning-analytics tools, onto third-party cloud environments. This migration has occurred alongside a parallel rise in reported cybersecurity incidents affecting the sector, including phishing campaigns, ransomware attacks, and credential theft, that have exposed both administrative and academic systems to serious disruption (Afolalu & Tsoeu, 2025; Ulven & Wangen, 2021). Unlike a conventional enterprise migrating a single business application to the cloud, HEIs migrating their core learning infrastructure inherit cloud-specific risk not for one system, but for an entire, increasingly interconnected ecosystem of adaptive, sensor- and data-rich learning technologies.

This paper adopts the term smart learning environments (SLEs) to describe this broader ecosystem, following Hwang's (2014) foundational characterization of SLEs as context-aware, adaptive, and, by design, cloud-dependent systems that sense learner context and adjust content, feedback, and guidance accordingly. An SLE is therefore not simply a cloud-hosted LMS; it is a layered set of interacting cloud services, adaptive content engines, learning-analytics dashboards, and, increasingly, AI-based tutoring and proctoring tools, each generating and consuming data about students in real time.

A clear gap emerges when the two relevant bodies of literature are placed side by side. Education-sector cybersecurity research documents the threat landscape facing HEIs in considerable empirical depth, but it rarely imports a structured, quantifiable risk-management methodology from the enterprise risk management (ERM) literature (Afolalu & Tsoeu, 2025; Ulven & Wangen, 2021). Conversely, the enterprise ERM and cloud-security-baseline literature offers mature, auditable governance

models, but these models were developed for general enterprise cloud deployments and were not built to anticipate the distinctive data types generated by SLEs, particularly learning-analytics-derived inferences and adaptive-system telemetry. This paper addresses that gap directly.

A second gap concerns measurement. Even where ERM is invoked, the education-sector literature seldom quantifies the risk reduction attributable to specific governance and control choices; Ulven and Wangen (2021) observe that empirical, quantified risk research specific to higher education remains comparatively scarce. Yet university leaders allocate resources in monetary terms and under uncertainty: which controls to fund first, how much residual loss is tolerable, and whether a proposed governance reform pays for itself. A framework that cannot answer these questions risks remaining a compliance document rather than a decision instrument.

The objectives of this paper are therefore to: (i) characterize the cloud-security threat landscape specific to smart learning environments, drawing on recent systematic reviews of higher-education cybersecurity; (ii) adapt enterprise ERM and cloud-security baselines, specifically the baseline requirements and strategic-integration arguments developed by Alam et al. (2024a, 2024b), to the SLE context; (iii) propose an institutional risk-management framework that maps education-sector threats onto enterprise baseline controls, classifies SLE-specific data types by risk, and embeds vendor risk management as a first-class concern; (iv) evaluate the framework with a hybrid quantitative design that combines analytic hierarchy process (AHP) prioritization, Monte Carlo loss simulation in the tradition of Factor Analysis of Information Risk (FAIR), service-dependency cascade simulation, and stepwise return-on-security-investment analysis; and (v) discuss the governance and adoption implications of implementing the framework within the decentralized institutional structures typical of higher education.

Four research questions organize the analysis. **RQ1:** Which framework components should institutions prioritize, and how stable is that ranking when expert judgments are uncertain? **RQ2:** By how much does the ERM-based framework change the distribution of annual cyber-loss for a representative SLE, including its tail? **RQ3:** How does the framework alter systemic exposure when SLE services depend on one another? **RQ4:** Which assumptions most strongly determine whether the framework's benefits materialize, and in what sequence should institutions adopt its components?

The paper makes three contributions. First, it adapts the enterprise cloud-security baseline of Alam et al. (2024a) and the strategic ERM argument of Alam et al. (2024b) to SLEs through a threat-to-control relevance matrix, a four-tier SLE data classification, a three-lines governance layer, and a dedicated vendor-risk component. Second, it introduces an integrated quantitative evaluation pipeline (Figure 2) whose stages can be re-run with institutional data. Third, it derives an evidence-informed adoption sequence and implementation roadmap. A caveat applies throughout: every quantitative input is an illustrative, literature-informed assumption rather than a measurement from a single institution (Section 4.8). The results should therefore be read as a demonstration of how the framework behaves and which assumptions matter most, not as forecasts for any particular university.

The remainder of the paper is organized as follows. Section 2 reviews the literature on SLEs, higher-education cyber threats, ERM baselines, learning-analytics ethics, and quantitative cyber-risk methods. Section 3 presents the conceptual framework. Section 4 describes the quantitative methods.

Section 5 reports and discusses the results. Section 6 considers institutional implications and limitations, and Section 7 concludes with recommendations.

2. Literature Review

2.1 Smart Learning Environments

Hwang (2014) provided the foundational definition of smart learning environments, characterizing them as context-aware, ubiquitous learning systems capable of sensing learners' real-world and online contexts and adapting content, guidance, and feedback accordingly. This definitional framework positions cloud infrastructure not as an optional deployment choice but as a structural precondition for SLE functionality: context sensing, adaptive content delivery, and real-time learning analytics all depend on persistent, scalable, cloud-based data processing. This distinction matters for security analysis, because it means that a security failure in the underlying cloud infrastructure is not a peripheral IT problem for an SLE, it is a failure of the pedagogical system itself.

The scope of what counts as an SLE continues to widen. Systematic evidence on artificial intelligence in higher education shows rapid growth in adaptive systems, intelligent tutoring, profiling, and automated assessment, together with limited critical reflection on the ethical and risk implications of deploying them (Zawacki-Richter et al., 2019). For security analysis the implication is practical: each new AI-enabled service adds a data-hungry, usually vendor-operated component to the institutional attack surface, and each is typically adopted by an academic unit rather than through a central risk process.

2.2 Cybersecurity Threats in Higher Education's Cloud Infrastructure

Recent systematic reviews converge on a consistent threat profile for HEIs. Afolalu and Tsoeu (2025) synthesized recent literature on HEI cybersecurity and identified phishing, ransomware, credential-based attacks, and the growing exposure created by institutional reliance on third-party cloud providers as dominant and recurring threat categories, while also highlighting limited breach-disclosure practices as a structural barrier to sector-wide learning from incidents. Ulven and Wangen (2021), applying a Comprehensive Literature Review model to twelve years of published research, similarly found a high level of cross-study agreement on core HEI cybersecurity risk areas, including vulnerable assets, insider threats, and gaps in institutional risk-management maturity, while noting that empirical, quantified risk research specific to higher education remains comparatively scarce.

Both reviews converge on a critical finding for this paper's argument: HEI cybersecurity research is threat-rich, it documents what goes wrong with considerable consistency, but comparatively framework-poor, offering few structured, quantifiable methodologies for prioritizing and managing the identified risks at an institutional level. This paper takes the seven threat categories that recur across these reviews, phishing, credential theft, ransomware and malware, insider threats, third-party or vendor breach, denial of service, and misconfigured cloud storage, as the threat taxonomy for all subsequent analysis.

2.3 Enterprise Risk Management and Cloud Security Baselines

The enterprise risk management literature offers the structured methodology that the education-sector literature lacks. Alam et al. (2024a) developed a set of baseline security requirements for cloud computing situated explicitly within an enterprise risk

management framework, specifying controls across access control, data classification, encryption, vendor risk management, and incident response for general enterprise cloud deployment. Complementing this technical baseline, Alam et al. (2024b) examined the strategic integration of ERM into organizational decision-making, arguing that successful ERM implementation depends on sustained leadership commitment, the cultivation of a risk-aware organizational culture, and the formal embedding of risk assessment into institutional planning processes, rather than treating risk management as a discrete, isolated compliance exercise.

Together, these two papers supply both the technical control baseline (Sections 3.1 and 3.2 of this paper) and the governance model (Section 3.3) around which the proposed SLE framework is built. They sit within a wider family of risk standards and guidance: the COSO ERM framework (Committee of Sponsoring Organizations of the Treadway Commission [COSO], 2017) integrates risk with strategy and performance; ISO 31000 (International Organization for Standardization [ISO], 2018) prescribes principles and a continuous process; and NIST SP 800-30 (National Institute of Standards and Technology [NIST], 2012) describes how to conduct cyber-risk assessments. The 2024 revision of the NIST Cybersecurity Framework added an explicit Govern function, reflecting the same shift from technical to organizational ownership of cyber risk (NIST, 2024).

2.4 Data Privacy and Ethics in Educational Analytics

A third literature establishes why the enterprise baseline, however mature, functions as a floor rather than a ceiling once student data enters the picture. Prinsloo and Slade (2013), evaluating the policy frameworks of two higher education institutions, found that institutional data-governance policy had not kept pace with the expanding analytic capabilities of learning-analytics systems, leaving significant gaps between what institutions were technically capable of inferring about students and what their existing privacy policies anticipated or authorized.

Slade and Prinsloo (2013) extended this concern into a broader ethical framework, arguing that the ethical stakes of learning analytics extend beyond conventional data-security concerns to questions of consent, data ownership, and the moral status of algorithmically generated inferences about learners. This literature establishes that SLE-specific data, particularly analytics-derived inferences, carries governance obligations that exceed what a generic enterprise cloud-security baseline was designed to address.

2.5 Quantitative Methods for Cyber-Risk Decision Support

Four families of quantitative methods are relevant to evaluating an ERM-based framework. The first is probabilistic loss modeling. The FAIR approach decomposes risk into loss-event frequency and loss magnitude and expresses both as distributions rather than point estimates (Freund & Jones, 2015). Combined with Monte Carlo simulation (Metropolis & Ulam, 1949), it yields a full distribution of annual loss from which expected loss, value-at-risk, and loss-exceedance curves can be read; Hubbard and Seiersen (2016) argue that such curves communicate risk to executives more honestly than ordinal heat maps because they make tolerance thresholds explicit.

The second family is multi-criteria prioritization. The analytic hierarchy process (Saaty, 1980, 2008) derives priority weights from pairwise judgments and provides a consistency ratio that flags incoherent judgment sets, which makes it suitable when decision criteria are heterogeneous, such as regulatory

exposure, confidentiality, continuity, and financial loss. The third family concerns interdependence. Cascading compromise across connected services can be represented with the independent cascade model, originally formalized for diffusion on networks (Kempe et al., 2003), in which each newly affected node has a single probabilistic chance to affect each of its neighbors.

The fourth family addresses economics. Return on security investment (ROSI) compares the reduction in expected loss with the cost of a control (Sonnenreich et al., 2006), and Gordon and Loeb (2002) showed that, for a broad class of breach-probability functions, the optimal investment in protecting an information set is never more than about 37 percent (1/e) of the expected loss it faces, so marginal returns diminish quickly. Finally, because every input to such models is uncertain, sensitivity analysis is indispensable; one-at-a-time designs are transparent but cannot detect interactions, a limitation emphasized in the global sensitivity-analysis literature (Saltelli et al., 2008).

To our knowledge, these methods have rarely been combined in an SLE-specific, ERM-grounded evaluation. The education-sector literature supplies threats but few models; the enterprise literature supplies models but not SLE data types or the decentralized procurement structure of universities. The present study integrates them, which is the research gap it fills.

3. Conceptual Framework

3.1 Threat-to-Control Mapping

The first component of the proposed framework maps the HEI-specific threats identified in Section 2.2, phishing, credential theft, ransomware and malware, insider threats, third-party or vendor breach, denial-of-service, and misconfigured cloud storage, onto the enterprise baseline controls specified by Alam et al. (2024a): access control, data encryption, data classification, vendor due diligence, incident response, and security-awareness training. Table 1 records the relevance of each control to each threat on a 0 to 3 scale, where 0 indicates no material relevance and 3 indicates a primary control. The matrix is the paper's own mapping and is used later, rescaled to the unit interval, as the coverage input of the loss model.

Threat	AC	EN	DC	VD	IR	AT
Phishing	3	1	0	0	2	3
Credential theft	3	1	0	0	2	2
Ransomware / malware	2	1	1	1	3	2
Insider threat	3	2	2	0	2	3
Vendor / third-party breach	1	2	2	3	2	0
Denial of service	0	0	0	1	3	0
Misconfigured storage	1	3	3	1	1	1

Table 1 Relevance of baseline controls to SLE threats (0 = none, 3 = primary control). AC = access control, EN = encryption, DC = data classification, VD = vendor due diligence, IR = incident response, AT = awareness training.

Two patterns are notable in Table 1. First, no single control provides strong coverage across all threat categories; phishing and credential theft are best addressed by access control and awareness training, while ransomware and misconfigured storage depend more heavily on encryption, data classification, and incident response. This reinforces the ERM principle that a defensible security posture requires layered, complementary controls rather than reliance on any single safeguard. Second, vendor due diligence is the only control category with primary

relevance to third-party or vendor breach, the threat category most directly implicated by SLE cloud migration, underscoring why vendor risk management is treated as a dedicated framework component in Section 3.4 rather than folded into general access control practice.

3.2 Risk Classification for SLE Data Types

General enterprise cloud-security baselines typically specify a generic data-classification requirement without anticipating the specific data types an SLE generates. This paper extends that generic requirement into an SLE-specific classification scheme comprising four tiers: directory information; protected educational records falling within FERPA or GDPR scope; learning-analytics-derived inferences, such as risk scores and behavioral or performance profiles; and adaptive-system telemetry, including clickstream data and, where applicable, biometric or sensor data collected by adaptive or proctoring tools. Table 2 summarizes these tiers.

SLE Data Tier	Illustrative Data Examples	Sensitivity / Governance Notes
Directory Information	Name, institutional email, enrollment status, published course roster data	Lowest, generally low-risk if disclosed
Protected Educational Records	Grades, transcripts, disciplinary records, financial aid data	High; long-standing FERPA/GDPR scope
Learning-Analytics-Derived Inferences	Risk scores, engagement predictions, behavioral or performance profiles	High and largely novel; few enterprise baselines anticipate derived/inferred data
Adaptive-System Telemetry	Clickstream logs, time-on-task data, and, where applicable, biometric or sensor data from proctoring and adaptive tools	Highest and most novel; minimal existing governance coverage

Table 2 SLE data classification scheme. Tier definitions synthesized from Alam et al. (2024a) baseline data-classification requirements, extended using Prinsloo and Slade (2013) and Slade and Prinsloo (2013).

Protected educational records are highly sensitive but comparatively well governed, reflecting decades of FERPA and GDPR-driven policy development. Learning-analytics-derived inferences and adaptive-system telemetry, by contrast, combine high sensitivity with a substantial governance gap, precisely the combination that Prinsloo and Slade (2013) identified as outpacing institutional policy capacity. These two tiers constitute genuinely novel risk categories that general enterprise baselines do not anticipate, and they are the reason the framework treats data classification as an SLE-specific control rather than importing the generic requirement unchanged.

3.3 Governance Layer

The third framework component applies Alam et al.'s (2024b) strategic-integration argument directly to the SLE context: institutional cybersecurity should be governed as a strategic risk-management function, owned at the executive or board level, formally budgeted, and subject to regular audit, rather than delegated entirely to IT departments. To make this ownership operational, the framework adopts the three-lines structure of the Institute of Internal Auditors (2020), shown in Figure 1. Executive leadership sets risk appetite and funds the program; first-line IT and academic units operate controls; a second line of risk, privacy, and compliance functions maintains the data-classification scheme and vendor risk register; and internal audit provides independent assurance.




Figure 1 Proposed ERM-based governance architecture for smart learning environments, combining the three lines model, the vendor gate, SLE data tiers, and a continuous feedback loop.

The architecture in Figure 1 is deliberately cyclical rather than linear. The continuous monitoring, incident response, and audit loop feeds findings back to the governance layer, ensuring that lessons learned from incidents or vendor audits inform subsequent risk classification and control decisions, consistent with the continuous-improvement orientation of ISO 31000 risk-management guidance and the COSO ERM framework. The governance layer also performs a function that is easy to overlook: in decentralized universities it determines how consistently controls are actually applied across colleges, departments, and research centers. Section 4.3 operationalizes this idea as a governance-consistency factor.

3.4 Vendor and Third-Party Risk Management

HEIs increasingly outsource core learning infrastructure to third-party vendors, cloud LMS providers, learning-analytics platforms, AI tutoring and generative AI tools, and, in many institutions, online proctoring or biometric verification tools, and in doing so, they inherit each vendor's security posture as their own institutional risk. This component draws directly on the vendor and cloud-provider risk requirement specified in Alam et al. (2024a) and extends it into a gate that precedes procurement rather than a review that follows it.

In practice the gate has four elements: a pre-procurement risk assessment scored by data-access scope and service criticality; contractual clauses covering breach notification, sub-processor disclosure, audit rights, and data return or deletion at exit; a vendor risk register that feeds the institutional risk register; and periodic reassessment triggered by changes in the vendor's data scope or ownership. Vendors that touch learning-analytics inferences or adaptive-system telemetry, the two highest-governance-gap data tiers, receive the most intensive review. Section 5.6 quantifies why this emphasis matters by modeling how a compromise in one vendor-operated service can propagate to the rest of the SLE.

3.5 Operating Model and Decision Rules

The framework is designed to be run as an annual cycle with quarterly checkpoints. In the inventory step, the second line

maintains a register of SLE services, their vendors, and the data tiers each processes. In the assessment step, services are scored against the threat-to-control matrix and the vendor gate. In the quantification step, the risk office elicits calibrated ranges for the loss-event frequency and magnitude of the principal risks and runs the loss model of Section 4.3. In the decision step, the executive risk committee compares the resulting loss-exceedance curve with the stated risk appetite and funds components in the sequence indicated by the stepwise analysis. In the monitoring step, key risk indicators (Table 3) are reported quarterly, and incidents and audit findings update the model inputs.

Two decision rules make the cycle operational. A *tolerance rule* states a maximum acceptable probability of exceeding a given annual loss, for example 5 percent for a loss of \$5 million, and triggers escalation whenever the loss-exceedance curve crosses it. A *sequencing rule* funds components in descending order of marginal benefit-cost ratio subject to the available budget, with the AHP ranking used as a check that regulatory and continuity considerations are not neglected. Section 5.4 shows why both rules are needed: a cost-blind ranking and a benefit-blind ranking each leave avoidable loss on the table when budgets bind.

Component	Example key risk indicator	Owner
Governance layer	Named executive risk owner; share of SLE services with an assigned risk owner; board review of loss-exceedance curve	Board
Access control	Multi-factor coverage of faculty, staff, and privileged accounts; share of accounts with excess privilege	1st line
Encryption	Share of SLE data stores encrypted at rest and in transit; key-rotation compliance	1st line
Data classification	Share of SLE data assets assigned to a tier; count of unclassified analytics datasets	2nd line
Vendor due diligence	Share of vendors assessed before contract; open high-risk findings; contracts with breach-notification clauses	2nd line
Incident response	Mean time to detect and contain; exercise completion; time to regulatory notification	1st / 2nd line
Awareness training	Training completion; simulated-phishing report rate; repeat-click rate	1st line
Risk quantification	Share of principal risks with loss-distribution estimates; back-testing of estimates against incidents	Audit

Table 3 Illustrative key risk indicators for monitoring each framework component, with the line of defense that owns each indicator.

The framework does not replace established standards but instantiates them for SLEs. Table 4 maps each component to the component structure of COSO ERM (2017), the process of ISO 31000 (2018), and the functions of the NIST Cybersecurity Framework 2.0 (2024), so that institutions already aligned with one of these can adopt the framework without building a parallel structure.

Component	COSO ERM (2017)	ISO 31000 (2018)	NIST CSF 2.0
Governance layer	Governance and culture	Leadership and commitment	Govern (oversight)
Data classification	Performance (identify, assess)	Risk assessment	Identify; Protect (data security)
Access control	Performance (respond)	Risk treatment	Protect (identity and access)

Component	COSO ERM (2017)	ISO 31000 (2018)	NIST CSF 2.0
Encryption	Performance (respond)	Risk treatment	Protect (data security)
Vendor due diligence	Performance (respond)	Risk treatment	Govern (supply-chain risk)
Incident response	Performance (respond)	Risk treatment	Detect; Respond; Recover
Awareness training	Governance and culture	Communication and consultation	Protect (awareness and training)
Risk quantification	Information, communication, reporting	Risk analysis	Govern (risk management strategy)
Audit and feedback loop	Review and revision	Monitoring and review	Govern (oversight); Identify (improvement)

Table 4 Crosswalk of framework components to COSO ERM, ISO 31000, and NIST CSF 2.0.

3.6 Propositions for Evaluation

The framework embodies five propositions that the quantitative analysis in Sections 4 and 5 is designed to probe. **P1 (consistency):** in decentralized institutions, the consistency with which controls are applied across units is a first-order determinant of loss, comparable in effect to the choice of controls themselves. **P2 (vendor amplification):** weak assurance of vendor-operated services amplifies systemic exposure beyond what their individual compromise probabilities imply. **P3 (prioritization):** ranking components by multi-criteria importance gives a different order from ranking them by cost-effectiveness, and the two serve different decisions. **P4 (robustness):** the expected-loss reduction attributable to the framework is robust to moderate uncertainty in individual parameters. **P5 (tail concentration):** as expected loss falls, the residual loss becomes concentrated in rare, severe years. Because the model is built from the framework's own mechanisms, support for P1 and P2 in particular shows that those mechanisms behave coherently, not that they have been validated empirically; Section 6.1 returns to this distinction.

4. Materials and Methods

4.1 Analytical Design

The evaluation follows the six-stage pipeline in Figure 2. Stages 1 and 2 are conceptual and were presented in Sections 2 and 3. Stages 3 to 6 are quantitative: AHP prioritization of framework components, Monte Carlo loss quantification, cascade and sensitivity analysis, and decision support. The analysis is performed for a representative mid-to-large HEI that operates a cloud-hosted SLE comprising twelve interdependent services. Three scenarios are compared throughout. The *inherent* scenario assumes no security controls and serves as a reference point. The *baseline* scenario represents typical decentralized governance, in which controls exist but are applied unevenly. The *ERM* scenario represents the target state of the proposed framework.




Figure 2 Analytical design: six-stage pipeline from evidence synthesis to decision support.

The framework has seven components: the six baseline control domains of Alam et al. (2024a), access control (AC), encryption (EN), data classification (DC), vendor due diligence (VD), incident response (IR), and awareness training (AT), plus the governance layer (GOV) of Section 3.3. Treating governance as a component alongside the technical controls is a modeling choice that allows its contribution to be compared with, and isolated from, theirs.

4.2 Analytic Hierarchy Process

The AHP evaluates the seven components against four criteria: regulatory and privacy exposure (C1), data confidentiality impact (C2), operational continuity of teaching and learning (C3), and financial and reputational loss (C4). Criteria weights are obtained from the pairwise comparison matrix A in Table 5, using Saaty's 1 to 9 scale, as the principal eigenvector:

$$A w = \lambda_{\max} w, \quad CR = (\lambda_{\max} - n) / [(n - 1) RI] \quad (1)$$

Criterion	C1	C2	C3	C4	Weight
C1 Regulatory & privacy exposure	1	1	2	2	0.326
C2 Data confidentiality impact	1	1	2	3	0.363
C3 Operational continuity	1/2	1/2	1	1	0.163
C4 Financial & reputational loss	1/2	1/3	1	1	0.148

Table 5 AHP pairwise comparison matrix of criteria and derived weights ($\lambda_{\max} = 4.021$; CI = 0.0069; CR = 0.0076).

where RI is Saaty's random consistency index (0.90 for $n = 4$) and a consistency ratio below 0.10 is conventionally accepted (Saaty, 1980). The judgments in Table 5 are the authors' illustrative baseline judgments, informed by the threat literature reviewed above; in an institutional application they would be replaced by the geometric-mean aggregate of a panel of risk, privacy, and IT leaders. Components are then scored on each criterion using the five-level rating (absolute measurement) mode of the AHP, in which levels 1 to 5 map to idealized

intensity priorities v of 0.08, 0.19, 0.36, 0.62, and 1.00. The global priority of component k is

$$P_k = \sum_c w_c v(r_{kc}) \quad (2)$$

where r_{kc} is the rating of component k on criterion c (Table 6). Because pairwise judgments and ratings are subjective, the ranking is stress-tested by a stochastic robustness analysis: in each of 10,000 replications, every pairwise judgment is shifted by one step on the Saaty scale with probability 0.25 in each direction, every rating is shifted by one level with probability 0.20 in each direction, and the priorities are recomputed. The share of replications in which each component attains each rank (a rank acceptability index) shows which parts of the ordering can be trusted.

Component	C1	C2	C3	C4	Priority	Rank
Access control	4	5	3	4	0.173	1
Encryption	4	5	2	3	0.157	3
Data classification	5	4	2	3	0.153	4
Vendor due diligence	4	4	3	4	0.139	5
Incident response	3	3	5	5	0.135	6
Awareness training	3	3	3	3	0.087	7
Governance layer	5	3	4	4	0.157	2

Table 6 Component ratings on each criterion (1 = very low to 5 = very high contribution) and resulting global AHP priority and rank.

4.3 Monte Carlo Loss Model

Annual loss is simulated in the FAIR tradition as the product of loss-event frequency and loss magnitude for each of the seven threats. For threat i , the inherent number of loss events per year is Poisson-distributed with a rate λ_i that is itself uncertain (lognormal, with log-standard deviation 0.35 around the median in Table 7). Each event has a lognormal severity X_{ij} . Controls act on this inherent profile through two channels, a thinning of events with probability f_i and a proportional reduction m_i in the severity of the events that remain, so that annual loss is

$$A_y = \sum_{i=1}^7 (1 - m_i) S_i \quad (3)$$

where S_i is the sum of the severities of the loss events retained for threat i . The frequency and magnitude reductions combine the contributions of the six control domains multiplicatively, which reflects the layered-defense assumption that independent controls each remove a share of what remains:

$$f_i = 1 - \prod_c [1 - a \varphi_c K_{ic} q_c g] \quad (4)$$

$$m_i = 1 - \prod_c [1 - a (1 - \varphi_c) K_{ic} q_c g] \quad (5)$$

Here K_{ic} is the relevance in Table 1 divided by three, q_c is the strength with which control c is implemented where it is applied, g is the governance-consistency factor (the share of institutional units and assets to which controls are in fact applied), $a = 0.80$ is the maximum share of a loss dimension that a perfect, universally applied control could remove, and φ_c is the share of control c 's effect that operates on frequency rather than magnitude. Preventive controls such as access control and awareness training have high φ , whereas encryption and incident response mainly limit impact (Table 8). The strength q_c and consistency g are drawn from beta distributions around the scenario means in Table 8, with concentration parameters 12 and 10 respectively.

Threat	Inherent events per year	Median loss per event (\$k)	Severity σ (log)
Phishing	6.0	35	1.1

Threat	Inherent events per year	Median loss per event (\$k)	Severity σ (log)
Credential theft	3.5	55	1.2
Ransomware / malware	1.5	600	1.3
Insider threat	2.0	120	1.2
Vendor / third-party breach	2.0	450	1.3
Denial of service	3.0	60	0.9
Misconfigured storage	3.2	220	1.2

Table 7 Threat parameters of the loss model (illustrative assumptions; median event rate, lognormal severity).

Component	Baseline mean	ERM mean	Freq. share ϕ	Cost (\$k/yr)
Access control	0.55	0.80	0.85	220
Encryption	0.50	0.80	0.15	180
Data classification	0.30	0.65	0.30	210
Vendor due diligence	0.25	0.65	0.60	160
Incident response	0.35	0.70	0.20	300
Awareness training	0.30	0.60	0.90	90
Governance (g)	0.60	0.85	n/a	250

Table 8 Control parameters of the loss model (illustrative assumptions): mean implementation strength q , governance consistency g , frequency share ϕ , and annualized cost of reaching ERM strength.

The parameters follow the qualitative pattern of the literature rather than any single dataset. Event rates are highest for phishing and credential-based attacks, which the reviewed studies report most often, and lowest for ransomware and vendor breaches, which are rarer but carry the heaviest tails; lognormal severity is a common practical choice for positive, right-skewed loss magnitudes (Hubbard & Seiersen, 2016). Baseline control strengths are lowest for data classification and vendor due diligence, consistent with reported gaps in institutional risk-management maturity (Ulven & Wangen, 2021), and highest for access control and encryption, which are better established in IT practice. The ERM strengths represent a deliberately attainable rather than perfect target state, and the cost figures are placeholders for staffing, tooling, and external services.

Each scenario is simulated with 100,000 trials. The same underlying random numbers are used for all scenarios (common random numbers), so differences between scenarios reflect the controls rather than sampling noise. Reported outputs are the mean annual loss (the annualized loss expectancy, ALE), the median, the 95th and 99th percentiles (value-at-risk, VaR), the conditional value-at-risk (CVaR, the mean loss in the worst 5 percent of years), exceedance probabilities at \$2 million and \$5 million, and the full loss-exceedance curve.

4.4 Service-Dependency Cascade Model

Because SLE services depend on one another, a single compromised service can trigger a chain of further compromises. The twelve services in Figure 12 (identity provider, core LMS, student information system, learning-analytics platform, AI tutoring service, online proctoring, cloud object storage, video conferencing, content delivery, e-mail and messaging, payment gateway, and mobile learning back end) are connected by 29 directed dependency edges. Each service has an annual probability of independent entry compromise, for example through phished credentials or a vendor-side fault. Once compromised, a service gets one chance to compromise each service that depends on it, with probability

$$\Pr(i \rightarrow j) = p_{ij} (1 - h_j) \quad (6)$$

where p_{ij} is the unmitigated propagation probability and h_j is the hardening of the target service through segmentation, least-privilege access, and vendor assurance. This is the independent cascade model (Kempe et al., 2003) applied to a security setting. Hardening is higher in the ERM scenario (0.60 for institution-operated and 0.55 for vendor-operated services) than at baseline (0.20 and 0.05), reflecting the weaker assurance typically available for third-party services in decentralized procurement, and entry probabilities scale in proportion to the remaining unhardened share. Each scenario is simulated 50,000 times. A cascade affecting four or more of the twelve services in a year, one third of the environment, is defined as a systemic event.

4.5 Stepwise Cost-Benefit Analysis

To translate the results into an adoption sequence, the components are upgraded from baseline to ERM strength one at a time in descending AHP priority, and the Monte Carlo model is re-run after each step. The marginal reduction in expected loss at step k , ΔALE_k , is compared with the annualized cost C_k of that component (Table 8) through

$$ROSI_k = (\Delta ALE_k - C_k) / C_k \quad (7)$$

(Sonnenreich et al., 2006). Costs are annualized, illustrative placeholders for staffing, tooling, and external services. Because the combined effect of controls is not additive, the marginal benefit of a component depends on what has already been implemented; the ladder therefore reports the benefit of each component in its policy-relevant position in the sequence rather than as a stand-alone effect.

Because the AHP ranks components by multi-criteria importance rather than by cost-effectiveness, two alternative sequences are also evaluated. The *benefit-cost greedy* sequence upgrades, at each step, the component with the highest marginal reduction in expected loss per dollar given what has already been implemented, re-evaluating all remaining components after every step. The *cost-ascending* sequence upgrades the cheapest components first and represents the common practice of starting with inexpensive, visible measures. All three sequences end in the same fully upgraded state, so they differ only in the path, and they are compared at three annual budget caps.

4.6 Sensitivity Analysis

Ten parameter groups are varied one at a time by 20 percent around their nominal values: inherent event rates, loss magnitude, severity dispersion, ERM control strength, ERM governance consistency, the control-effectiveness ceiling, the prevention-versus-impact split (shifted by 0.15 in either direction), the threat-control relevance weights, the ransomware event rate, and the vendor-breach loss magnitude. Parameters bounded by one are capped at 0.98. For each, the full model is re-run with 100,000 trials and the resulting ERM expected annual loss, and the expected-loss reduction relative to baseline, are recorded and displayed as a tornado chart.

4.7 Maturity Trajectory Model

Institutional capability rarely improves linearly: effort is front-loaded on mobilization, accelerates once practices become routine, and saturates near the target. Following the S-shaped adoption curves familiar from diffusion research (Rogers, 2003), the maturity of dimension d (leadership ownership, risk quantification, vendor governance, data classification, incident-response readiness, and regulatory alignment), scored 1 to 5, follows a logistic path re-based to start at its initial value:

$$M_d(t) = M_{d0} + (M_d^* - M_{d0}) [S_d(t) - S_d(0)] / [1 - S_d(0)] \quad (8)$$

$$S_d(t) = 1 / \{1 + \exp[-k_d(t - \tau_d)]\} \quad (9)$$

where t is months since launch, τ_d is the inflection point, and k_d the steepness. Inflection points are earliest for leadership ownership, which depends on a governance decision, and latest for vendor governance and risk quantification, which depend on contract cycles and data accumulation. Uncertainty is propagated by 2,000 replications that perturb inflection points and steepness lognormally (log-standard deviations 0.25 and 0.20) and targets normally (standard deviation 0.25 maturity points).

4.8 Assumptions, Reproducibility, and Validity

All model inputs are illustrative assumptions chosen to be of a plausible order of magnitude for a mid-to-large institution; none is estimated from institutional incident data. This is consistent with the conceptual character of the paper, but it limits interpretation. Absolute dollar values should not be read as predictions, whereas relative comparisons between scenarios, the ordering of components, and the identification of influential parameters are the substantive outputs. Every parameter is reported in Table 1 to Table 8 so that the analysis can be reproduced and re-run with local data. With 100,000 trials the standard error of the mean annual loss is below \$20,000 in every scenario, so Monte Carlo noise is small relative to the differences reported in Section 5.

Three verification checks support the implementation. First, with all control strengths set to zero the model reproduces the inherent loss exactly (a difference of \$0). Second, expected annual loss declines monotonically as ERM control strengths are scaled to 0, 0.5, 1.0, and 1.25 times their nominal values (\$7.63M, \$3.36M, \$1.36M, and \$0.90M). Third, estimates are stable with respect to the number of trials (standard errors at 100,000 trials are about \$12k and \$5k): from 5,000 to 100,000 trials the estimated ERM-versus-baseline reduction varies by only 0.3 percentage points while the standard errors shrink in proportion to the inverse square root of the number of trials.

Construct validity rests on the faithful translation of the framework into model parameters: the relevance matrix is the paper's own mapping, the control domains are those of the cited baseline, and the governance layer enters the model through the mechanism described in Section 3.3. External validity is limited by the single representative institution; Section 6.4 discusses how an empirical audit could address this.

5. Results and Discussion

5.1 Component Priorities and Rank Robustness

The criteria weights derived from Table 5 are 0.326 for regulatory and privacy exposure, 0.363 for data confidentiality impact, 0.163 for operational continuity, and 0.148 for financial and reputational loss. The principal eigenvalue is 4.021, giving a consistency index of 0.0069 and a consistency ratio of 0.0076, far below the 0.10 threshold, so the judgments are coherent. The weighting places data confidentiality and regulatory exposure first, which is consistent with the paper's premise that SLE data, rather than server availability alone, is the central asset at stake.

Figure 3 shows the resulting global priorities. Access control ranks first (0.173), followed closely by the governance layer (0.157), encryption (0.157), and data classification (0.153); vendor due diligence and incident response follow (0.139 and 0.135), and awareness training is last (0.087). The five middle components differ by only about two percentage points of global priority, so the point ranking should be read with caution. That the governance layer scores alongside the technical controls is itself informative: it ranks second even though it is the only component that does not directly stop an attack.




Figure 3 Pareto chart of global AHP priorities for the seven framework components (bars) and their cumulative share (line).

The robustness analysis in Figure 4 sharpens this reading. Access control attains the first rank in 43 percent of the 10,000 perturbed judgment sets and a top-two rank in 66 percent, whereas awareness training is ranked last in 88 percent. The ranks of the governance layer, encryption, data classification, vendor due diligence, and incident response, by contrast, are spread across positions two to six; the governance layer, for example, takes ranks 1 to 3 with probabilities of 18, 21, and 21 percent. The practical message is twofold. Institutions can be reasonably confident that identity and access hardening belongs at the front of the sequence and that awareness training, while worthwhile, should not be the first investment. For the middle five components, however, priority should be settled by local conditions such as cost, existing capability, and procurement calendars rather than by the generic ranking.




Figure 4 Rank acceptability of framework components across 10,000 perturbed AHP judgment sets (cell labels in percent).

5.2 Annual Loss Distributions

Table 9 summarizes the simulated distributions of annual loss. With no controls the mean annual loss is \$7.63 million. Typical decentralized governance reduces it to \$4.15 million (46% below inherent), and the ERM scenario to \$1.36 million, which is 67.4% below baseline and 82% below the inherent level. Tail measures fall by a similar proportion: the 95th percentile declines from \$10.46 million to \$3.84 million (63%), and the conditional value-at-risk from \$15.59 million to \$6.17 million (60%). Reductions in the tail are slightly smaller than in the mean, a familiar feature of heavy-tailed losses: severe years are dominated by a few large events, and those are not eliminated, only made rarer and smaller.

Metric	Inherent	Baseline	ERM	Change vs baseline
Mean (ALE)	7.63	4.15	1.36	-67%
Median	6.19	3.17	0.93	-71%
VaR 95%	17.72	10.46	3.84	-63%
VaR 99%	29.39	18.05	7.25	-60%
CVaR 95%	25.47	15.59	6.17	-60%
P(loss > \$2M)	96.1%	74.1%	18.3%	-55.8 pp
P(loss > \$5M)	63.4%	26.2%	2.6%	-23.5 pp

Table 9 Simulated annual loss under the three scenarios (USD million unless stated; 100,000 trials; pp = percentage points).




Figure 5 Distribution of simulated annual loss under the baseline and ERM scenarios (violins, log scale; markers show median, mean, and 95th percentile).

Figure 5 shows the shape of the two distributions. Both are strongly right-skewed, with the mean above the median, but the skew is more pronounced under the ERM framework (mean-to-median ratio of 1.46 against 1.31 at baseline), because controls remove many moderate losses while leaving the rare severe events comparatively less affected. The entire distribution moves downward and narrows, which matters for budgeting as well as for tolerance: annual loss becomes more predictable as well as smaller.

The loss-exceedance curves in Figure 6 make the board-level implication explicit. Against an illustrative tolerance of \$5 million in annual loss, the baseline institution has a 26.2% probability of breaching it in any year, roughly one year in four, whereas under the ERM framework the probability is 2.6%. The probability of exceeding \$2 million falls from 74% to 18%. The ERM curve lies below the baseline curve across the entire range of thresholds, which means the framework does not trade a lower average for a worse tail.




Figure 6 Loss-exceedance curves for annual loss under the inherent, baseline, and ERM scenarios (log-log scale) with an illustrative \$5 million tolerance.

Although losses in the tail fall in absolute terms, they account for a larger share of what remains. The worst 5 percent of years contribute 19% of baseline expected loss but 23% of expected loss under the ERM framework. Lowering the mean therefore does not remove the need for tail-focused measures such as incident-response retainers, financial reserves, or cyber insurance, which the framework treats as complements to, not substitutes for, control improvement.

Figure 7 and Table 10 decompose the result by threat. In absolute terms the largest expected losses at baseline come from vendor or third-party breach and ransomware, which together account for more than half of baseline expected loss even though their event frequencies are modest; their severity is high. These are also the threats for which the framework delivers the largest absolute reductions. In relative terms, insider threats and misconfigured storage show the steepest declines, because several controls with high relevance act on them simultaneously (access control, awareness training, encryption, and data classification). Denial of service improves least, because the mapping in Table 1 assigns it only one primary control, incident response, and a minor role for vendor due diligence; mitigating volumetric attacks relies largely on provider-side capacity, which this framework influences mainly through vendor selection.




Figure 7 Expected annual loss by threat under the inherent, baseline, and ERM scenarios (log scale; labels give the ERM change relative to baseline).

Threat	Events/yr baseline	Events/yr ERM	ALE baseline (\$k)	ALE ERM (\$k)
Phishing	4.18	2.00	217	73

Threat	Events/yr baseline	Events/yr ERM	ALE baseline (\$k)	ALE ERM (\$k)
Credential theft	2.55	1.38	235	89
Ransomware / malware	1.14	0.64	1190	376
Insider threat	1.34	0.59	234	55
Vendor / third-party breach	1.69	1.07	1233	381
Denial of service	3.01	2.64	231	138
Misconfigured storage	2.68	1.78	811	243
Total	16.60	10.09	4151	1355

Table 10 Expected loss events per year and expected annual loss (ALE) by threat under the baseline and ERM scenarios.

5.3 Stepwise Adoption and Return on Investment

Figure 8 traces the expected annual loss as components are upgraded in AHP order, and Table 11 reports the economics. The governance step delivers the largest single reduction (\$1.06M), more than the whole of the first technical step (\$0.29M), because raising the consistency with which every control is applied multiplies the effect of all of them. This is the quantitative counterpart of the argument of Alam et al. (2024b): leadership ownership and cross-unit consistency are not overhead but the mechanism through which technical controls achieve their potential coverage. Upgrading all seven components reduces expected loss by \$2.80M per year at an annualized cost of \$1.41M, a net benefit of \$1.39M and an overall ROSI of 98%.




Figure 8 Stepwise reduction in expected annual loss as components are upgraded in AHP priority order (AC = access control, GOV = governance layer, EN = encryption, DC = data classification, VD = vendor due diligence, IR = incident response, AT = awareness training).

Step	Component upgraded	ΔALE (\$k)	Cost (\$k/yr)	ROSI
1	Access control (AC)	286	220	30%
2	Governance layer (GOV)	1055	250	322%
3	Encryption (EN)	384	180	113%
4	Data classification (DC)	295	210	40%
5	Vendor due diligence (VD)	272	160	70%
6	Incident response (IR)	367	300	22%
7	Awareness training (AT)	136	90	51%
Total (all seven)		2796	1410	98%

Table 11 Stepwise adoption in AHP priority order: marginal reduction in expected annual loss, annualized cost, and return on security investment (ROSI).

Marginal returns differ markedly by step. The governance step has the highest ROSI (322%), whereas incident response, the most expensive component, has the lowest (22%), although it remains positive. Figure 9 plots cumulative benefit against cumulative cost: the benefit curve rises steeply at the governance step and flattens thereafter, so each subsequent dollar buys less risk reduction. This is the diminishing-returns pattern anticipated by Gordon and Loeb (2002). The total ERM program cost equals 34% of baseline expected loss, just below the one-in-e ceiling of about 37 percent that Gordon and Loeb derive for the optimal outlay, which suggests that the full program sits at the upper end of what is economically efficient and that institutions with tighter budgets should stop earlier in the sequence.




Figure 9 Cumulative loss reduction versus cumulative annualized cost across the seven adoption steps.

5.4 Budget-Constrained Adoption Strategies

The analysis in Section 5.3 followed the AHP order. Figure 10 compares it with the benefit-cost greedy and cost-ascending sequences, and Table 12 reports the expected annual loss attainable at three budget caps. The greedy sequence is the most effective at every cap. With \$0.50 million per year it reaches an expected loss of \$2.47M by funding GOV, AT, VD, compared with \$2.81M for the AHP order (AC, GOV) and \$3.24M for the cost-ascending order (AT, VD, EN). At \$1.00 million the corresponding figures are \$1.73M for the greedy order and \$2.13M for the AHP order.




Figure 10 Expected annual loss against cumulative annualized cost for three adoption sequences (AHP order, benefit-cost greedy, and cost-ascending); dotted lines mark the \$0.50 million and \$1.00 million budget caps.

Annual budget	AHP order	Benefit-cost greedy	Cost-ascending
\$500k	2.81 (AC, GOV)	2.47 (GOV, AT, VD)	3.24 (AT, VD, EN)
\$750k	2.43 (AC, GOV, EN)	2.14 (GOV, AT, VD, EN)	2.97 (AT, VD, EN, DC)
\$1000k	2.13 (AC, GOV, EN, DC)	1.73 (GOV, AT, VD, EN, IR)	2.77 (AT, VD, EN, DC, AC)

Table 12 Expected annual loss (USD million, bold) and components funded under each adoption sequence at three annual budget caps.

The greedy sequence is GOV → AT → VD → EN → IR → DC → AC. It starts with the governance layer, as the stepwise analysis would suggest, but then moves awareness training, which the AHP ranks last, to second place, because it is inexpensive and its benefit is largest before other controls have reduced the losses it addresses. Access control, which the AHP ranks first, enters last, because it is relatively costly and much of its benefit is already captured by the preceding components. This illustrates the order dependence noted in Section 4.5: the value of a component is not intrinsic but depends on what has been implemented before it. The cost-ascending order is the least effective until its final step, because it defers the governance layer and incident response, the two components with the largest marginal reductions.

5.5 Sensitivity of the Results

Figure 11 ranks the parameter groups by the swing they induce in the ERM expected annual loss, whose nominal value is \$1.36M. The four most influential groups are erm control strength (all controls) (\$0.91 to 1.98 million), erm governance consistency (\$0.97 to 1.85 million), threat-control relevance weights (\$1.10 to 1.98 million), and severity dispersion (log-sd) (\$1.04 to 1.89 million). Two of these, control strength and governance consistency, are quantities that the institution itself determines through implementation quality, which reinforces the view that the framework's value lies in execution rather than in the choice of controls on paper. A 20 percent shortfall in ERM control strength, for example, raises expected loss by roughly 46 percent.




Figure 11 Tornado chart: one-at-a-time sensitivity of ERM expected annual loss to 20% parameter changes (see Section 4.6 for the ranges used for bounded parameters).


Parameters that scale the inherent exposure, event rates and loss magnitude, move expected loss almost proportionally (\$1.08 to 1.64 million for a 20 percent change), as expected for a model that is linear in both. Severity dispersion matters more than its modest range suggests because the mean of a lognormal

distribution grows with the square of its log-standard deviation, so heavier tails raise the expected loss even when medians are unchanged. The prevention-versus-impact split has almost no effect on the expected value (\$1.31 to 1.36 million), which implies that conclusions about the framework do not hinge on how a control's benefit is apportioned between lowering frequency and limiting damage. Most importantly for the credibility of the headline result, the expected-loss reduction relative to baseline never falls below 52% in any of the one-at-a-time cases, so the qualitative finding that the framework substantially lowers expected loss is robust to the assumptions varied here, although interactions among parameters are not explored by this design.

5.6 Systemic Cascade Risk

Figure 12 shows the annual compromise probability of each service under the two governance regimes, and Table 13 summarizes the cascade metrics. At baseline the expected number of compromised services is 2.01 per year, and in 23.1% of years at least four of the twelve services are compromised, a systemic event. Under the ERM framework the corresponding figures are 0.68 and 3.8%. The probability that at least one service is compromised in a year also falls, from 64% to 38%, but this decline is smaller than the decline in systemic events, which indicates that the framework's main effect is to prevent localized incidents from becoming cascades rather than to eliminate incidents altogether.

(a) Baseline: decentralized governance (mean 2.01 services compromised/yr)



(b) ERM framework (mean 0.68 services compromised/yr)




Figure 12 Annual compromise probability of the twelve SLE services under (a) baseline and (b) ERM governance; arrow width indicates effective propagation probability.

Metric	Baseline	ERM
Mean services compromised per year	2.01	0.68
P(at least one service compromised)	63.8%	37.9%
P(four or more services): systemic event	23.1%	3.8%
P(six or more services)	11.1%	0.7%
Amplification factor	2.08	1.46

Table 13 Cascade metrics for the twelve-service SLE dependency network (50,000 trials per scenario).

The most exposed services at baseline are the learning-analytics platform (0.29), the ai tutoring service (0.25), and the cloud object storage (0.23). Exposure is driven less by entry probability than by position in the dependency graph: services that sit downstream of many others accumulate risk from every upstream compromise. The mean compromise probability of vendor-operated services falls from 0.17 to 0.05, and that of institution-operated services from 0.16 to 0.06; the larger proportional reduction for vendor-operated services reflects the low assurance assumed for them at baseline and is the quantitative argument for the vendor gate of Section 3.4. The amplification factor, the expected number of compromised services per initially compromised service, falls from 2.08 to 1.46.

Service	Operator	Baseline	ERM
Learning-analytics platform	Vendor	29.0%	8.2%
AI tutoring service	Vendor	25.1%	7.3%
Cloud object storage	Vendor	22.7%	6.5%
E-mail & messaging	Institution	21.3%	9.7%
Identity provider (SSO)	Institution	18.5%	8.4%
Core LMS	Institution	17.6%	5.5%
Online proctoring	Vendor	16.6%	5.7%
Video conferencing	Vendor	13.2%	4.3%
Student information system	Institution	11.1%	3.8%
Mobile learning backend	Institution	10.2%	3.9%
Payment gateway	Vendor	8.1%	2.4%
Content / CDN	Vendor	7.3%	2.4%

Table 14 Annual compromise probability by service, sorted by baseline exposure.




Figure 13 Distribution of cascade size (services compromised per year); the shaded region marks systemic events (four or more services).

Figure 13 shows the full distribution of cascade sizes. At baseline the distribution has a long right tail, with 11.1% of years affecting six or more services. Under the ERM framework that probability is 0.7%. For an institution whose teaching depends on the simultaneous availability of its identity provider, LMS, and student information system, the compression of this tail is arguably more valuable than the reduction in the mean, because multi-service outages are the events that interrupt teaching and trigger regulatory notification.

5.7 Maturity Trajectory and Implementation Roadmap

Figure 14 projects the maturity of six governance dimensions over 36 months. The composite index rises from 1.90 at launch

to 3.02 at 12 months, 4.18 at 24 months, and 4.37 at 36 months, with a 10th to 90th percentile band of 2.79 to 3.26 at 12 months, narrowing to 4.23 to 4.50 at 36 months. Leadership ownership matures fastest, from 2.0 to 4.2 within the first year, because it depends on a decision rather than on accumulated practice. Vendor governance and risk quantification are slowest (2.1 and 2.3 at 12 months), as they depend on contract renewal cycles and on the accumulation of incident and loss data. The model implies that institutions should expect visible governance gains within a year but should not expect quantified, vendor-integrated risk management before the second year.




Figure 14 Projected maturity trajectories of six governance dimensions under the logistic model, with the composite index and its 10th to 90th percentile band.

Figure 15 translates these results into a phased 36-month roadmap whose order follows the findings above. The governance charter and the appointment of an executive risk owner come first, because the governance layer has the largest marginal effect and is a precondition for consistent control application. Baseline assessment and data tiering run in parallel and produce the classification scheme of Section 3.2. Identity, access, and encryption hardening, the highest-priority technical components, follow, with the vendor risk program starting by month six so that its register is complete by month 19, before the next major procurement cycle. Incident response and monitoring uplift culminates in a full-scale exercise at month 25; the awareness and culture program runs throughout, since its value depends on repetition; and risk quantification and the audit loop begin in the second year, once enough data exist to replace illustrative parameters with institutional ones.




Figure 15 Phased 36-month implementation roadmap (diamonds mark milestones: charter ratified, data tiers complete, vendor register complete, incident-response exercise, first independent audit).

6. Institutional Implications and Limitations

6.1 Evaluation of the Propositions

Table 15 relates each proposition from Section 3.6 to the model evidence. All five are supported within the model, but the nature of the support differs. P1 and P2 follow in part from how the model is constructed, since governance consistency and vendor hardening enter it as explicit parameters, so they should be read as confirmations that the mechanisms matter quantitatively rather than as independent discoveries. P3, P4, and P5, by contrast, are emergent results: the divergence between the AHP and benefit-cost orders, the stability of the loss reduction under one-at-a-time variation, and the rising share of loss in the tail were not imposed on the model and could have come out differently.

Proposition	Model evidence	Assessment
P1 Consistency	Governance step gives the largest loss reduction (\$1.06M, 38% of the total); governance consistency swings ERM loss between \$0.97M and \$1.85M	Supported within model
P2 Vendor amplification	3 of the 3 most exposed services at baseline are vendor-operated; their mean compromise probability falls from 0.17 to 0.05; amplification falls from 2.08 to 1.46	Supported within model
P3 Prioritization	At a \$0.50M budget the greedy order reaches \$2.47M against \$2.81M for the AHP order; AHP ranks awareness training last, the greedy order funds it second	Supported
P4 Robustness	Loss reduction never falls below 52% in any one-at-a-time case	Supported; interactions untested
P5 Tail concentration	Worst 5% of years contribute 19% of baseline and 23% of ERM expected loss	Supported

Table 15 Evaluation of the five propositions against the model evidence.

6.2 Barriers to Adoption

Several institutional barriers complicate the direct transplantation of enterprise ERM models into higher education. Universities typically operate decentralized IT governance structures, in which individual colleges, departments, or even research centers may procure and manage their own cloud-based learning tools outside centralized IT oversight, a structural condition that stands in tension with the centralized control assumed by most enterprise ERM models (Alam et al., 2024b). The model results make the cost of this condition visible: the difference between the baseline and ERM scenarios is driven substantially by the governance-consistency factor, which is precisely the share of the institution to which controls actually apply.

Budget constraints compound this barrier: unlike large commercial enterprises, many HEIs lack dedicated risk-management budgets separate from general IT operating expenses, making executive risk ownership, as recommended in Section 3.3, a genuine organizational change rather than a reporting adjustment. Faculty and staff security-awareness gaps represent a further recurring barrier documented throughout the HEI cybersecurity literature (Afolalu & Tsoeu, 2025; Ulven & Wangen, 2021), and they interact directly with the threat-to-control mapping in Table 1, where security-awareness training is among the controls most consistently relevant across phishing, credential theft, and insider-threat categories.

Awareness is not only a matter of information. Research on technology acceptance shows that use of a technology depends on perceived performance benefits, effort, social influence, and the facilitating conditions provided by the organization (Venkatesh et al., 2003), and work on teachers' technology

change shows that knowledge, confidence, beliefs, and institutional culture intersect (Ertmer & Ottenbreit-Leftwich, 2010). Applied to security, this suggests that faculty adoption of multi-factor authentication or vendor-approval procedures will depend on whether these measures are perceived as supportive of teaching, are easy to use, and are visibly endorsed by academic leaders, which is a further argument for placing ownership at the executive level.

Regulatory complexity further shapes how the enterprise baseline must be adapted rather than adopted wholesale. FERPA, GDPR, and an expanding patchwork of U.S. state-level student-data-privacy statutes impose obligations that extend beyond, and in some cases constrain, the generic enterprise cloud-security baseline, particularly with respect to the learning-analytics-derived inferences and adaptive-system telemetry identified as high-governance-gap data types in Table 2. A framework that treats these regulatory regimes as an added compliance checklist, rather than as a structural input into data classification and vendor due diligence, risks reproducing the very governance gap that Prinsloo and Slade (2013) identified over a decade ago.

6.3 Positioning and Practical Implications

Positioned against existing HEI security guidance, much of which consists of relatively generic institutional IT security policy documents, the proposed framework is more prescriptive and risk-quantified in four respects. It specifies an explicit data-classification scheme for SLE-specific data types (Table 2); it operationalizes a threat-to-control relevance matrix rather than a generic control checklist (Table 1); it expresses risk as a loss distribution with explicit tolerance thresholds rather than as an ordinal rating (Figure 6); and it treats vendor risk exposure as a comparatively scored, prioritizable input to institutional risk registers rather than a binary procurement gate (Figure 12).

Three practical uses follow. First, the loss-exceedance curve is a natural board-reporting artifact: leaders can state a tolerance, read off the probability of breaching it, and observe how that probability changes with proposed investments. Second, the stepwise ladder gives budget committees a defensible sequence and a stopping rule, namely to proceed while the marginal ROSI is acceptable. Third, the cascade model gives enterprise architects a way to compare segmentation and vendor-assurance options by their effect on systemic events rather than on single-service compromise. This positioning is consistent with the broader argument advanced throughout this paper: that education-sector security practice benefits from importing the structured, quantifiable methodology of enterprise ERM, while enterprise ERM benefits from the SLE-specific data sensitivity insights developed within the learning-analytics ethics literature.

6.4 Limitations and Future Research

Several limitations qualify the findings. First, all inputs are illustrative assumptions; the dollar figures describe how the model behaves, not what any institution should expect. Second, the model treats threats as independent apart from the cascade analysis, ignoring correlation such as a single phishing campaign precipitating both credential theft and ransomware. Third, the multiplicative combination of controls assumes independence of their effects, which may overstate benefits where controls share failure modes. Fourth, the AHP judgments come from the authors rather than from a panel, and the one-at-a-time sensitivity design cannot reveal interactions among parameters. Fifth, the cascade model represents dependencies as static probabilistic edges and does not capture detection and containment dynamics once an intrusion is under way.

Future research should pursue an empirical audit of a real institution's cloud-security posture against the baseline proposed here, replacing the illustrative parameters used throughout Section 4 with institution-specific risk-register data, incident logs, and vendor security documentation. Calibrated expert elicitation could supply the pairwise judgments, and global sensitivity methods such as variance-based indices could quantify interactions (Saltelli et al., 2008). Such work would also allow the threat-to-control relevance matrix and the cascade structure to be validated and refined across institutional contexts that vary in size, decentralization, and existing security maturity.

7. Conclusion and Recommendations

This paper has argued that the cloud-specific cybersecurity and data-privacy risks documented extensively in the higher-education security literature have not yet been paired with a formal, enterprise-grade risk-management response, and it has proposed an ERM-based cloud security framework for smart learning environments that closes this gap. By mapping HEI-specific threats onto enterprise baseline controls, extending generic data classification into an SLE-specific four-tier scheme, embedding cybersecurity governance as a strategic, leadership-owned function, and treating vendor risk management as a first-class framework component, the proposed model offers institutions a more prescriptive and risk-quantified alternative to generic IT security policy.

The quantitative evaluation answers the four research questions. AHP prioritization places access control first and awareness training last with high stability, while the middle five components are close to tied (RQ1). Under illustrative parameters the ERM scenario lowers expected annual loss from \$4.15 million to \$1.36 million and the probability of exceeding a \$5 million loss from 26.2% to 2.6% (RQ2). It cuts the probability of a systemic cascade from 23.1% to 3.8% (RQ3). Control strength and governance consistency dominate output uncertainty, the governance layer offers the highest return and should lead the adoption sequence, and under binding budgets a benefit-cost ordering reaches a lower expected loss than the AHP ordering at every budget examined (RQ4).

Several practical recommendations follow directly from this analysis. First, institutions should treat the cloud-security baseline outlined in Sections 3.1 and 3.2 as a pre-migration requirement, evaluated before an LMS, analytics platform, or AI tutoring tool is adopted, rather than as a post-incident response. Second, institutions should assign explicit executive risk ownership for SLE cybersecurity, consistent with the governance layer proposed in Section 3.3 and the strategic-integration argument of Alam et al. (2024b), rather than leaving risk ownership implicitly distributed across decentralized IT units. Third, institutions should require documented, comparative vendor risk assessments before finalizing EdTech procurement contracts, with particular scrutiny applied to AI tutoring, learning-analytics, and proctoring tool vendors, whose position in the dependency network makes them the most exposed services. Fourth, institutional data-governance policy should be revisited specifically to address learning-analytics-derived inferences and adaptive-system telemetry, the two SLE data tiers this paper identifies as carrying the largest governance gap. Fifth, boards should adopt loss-exceedance reporting against an explicit tolerance and use stepwise return-on-investment analysis to sequence funding.

In sum, securing the smart campus is not primarily a technical challenge; it is a governance challenge, and the enterprise risk management literature offers higher education a structured, mature methodology for meeting it.

Conflict of Interest

The authors state that there is no conflict of interest.

References

- Afolalu, O., & Tsoeu, M. S. (2025). Cybersecurity in higher education institutions: A systematic review of emerging trends, challenges and solutions. *Future Internet*, 17(12), Article 575. <https://doi.org/10.3390/fi17120575>
- Alam, M. R. U., Shohel, A., & Alam, M. (2024a). Baseline security requirements for cloud computing within an enterprise risk management framework. *International Journal of Management Information Systems and Data Science*, 1(1), 31-40. <https://doi.org/10.62304/ijmisds.v1i1.115>
- Alam, M. R. U., Shohel, A., & Alam, M. (2024b). Integrating enterprise risk management (ERM): Strategies, challenges, and organizational success. *Global Mainstream Journal*, 1(2), 10-19. <https://doi.org/10.62304/ijbm.v1i2.130>
- Committee of Sponsoring Organizations of the Treadway Commission. (2017). *Enterprise risk management: Integrating with strategy and performance*. COSO.
- Ertmer, P. A., & Ottenbreit-Leftwich, A. T. (2010). Teacher technology change: How knowledge, confidence, beliefs, and culture intersect. *Journal of Research on Technology in Education*, 42(3), 255-284.
- Freund, J., & Jones, J. (2015). *Measuring and managing information risk: A FAIR approach*. Butterworth-Heinemann.
- Gordon, L. A., & Loeb, M. P. (2002). The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4), 438-457. <https://doi.org/10.1145/581271.581274>
- Hubbard, D. W., & Seiersen, R. (2016). *How to measure anything in cybersecurity risk*. Wiley.
- Hwang, G.-J. (2014). Definition, framework and research issues of smart learning environments: A context-aware ubiquitous learning perspective. *Smart Learning Environments*, 1, Article 4. <https://doi.org/10.1186/s40561-014-0004-5>
- Institute of Internal Auditors. (2020). *The IIA's three lines model: An update of the three lines of defense*. The Institute of Internal Auditors.
- International Organization for Standardization. (2018). *ISO 31000:2018 Risk management: Guidelines*. ISO.
- Kempe, D., Kleinberg, J., & Tardos, E. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 137-146). ACM. <https://doi.org/10.1145/956750.956769>
- Metropolis, N., & Ulam, S. (1949). The Monte Carlo method. *Journal of the American Statistical Association*, 44(247),

335-341.

<https://doi.org/10.1080/01621459.1949.10483310>

National Institute of Standards and Technology. (2012). Guide for conducting risk assessments (NIST Special Publication 800-30, Rev. 1). U.S. Department of Commerce.

<https://doi.org/10.6028/NIST.SP.800-30r1>

National Institute of Standards and Technology. (2024). The NIST Cybersecurity Framework (CSF) 2.0 (NIST CSWP 29). <https://doi.org/10.6028/NIST.CSWP.29>

Prinsloo, P., & Slade, S. (2013). An evaluation of policy frameworks for addressing ethical considerations in learning analytics. In *Proceedings of the 3rd International Conference on Learning Analytics and Knowledge (LAK '13)* (pp. 240-244). ACM.

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Saaty, T. L. (1980). *The analytic hierarchy process: Planning, priority setting, resource allocation*. McGraw-Hill.

Saaty, T. L. (2008). Decision making with the analytic hierarchy process. *International Journal of Services Sciences*, 1(1), 83-98.

<https://doi.org/10.1504/IJSSCI.2008.017590>

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). *Global sensitivity analysis: The primer*. Wiley.

Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist*, 57(10), 1510-1529.

<https://doi.org/10.1177/0002764213479366>

Sonnenreich, W., Albanese, J., & Stout, B. (2006). Return on security investment (ROSI): A practical quantitative model. *Journal of Research and Practice in Information Technology*, 38(1), 45-56.

Ulven, J. B., & Wangen, G. (2021). A systematic review of cybersecurity risks in higher education. *Future Internet*, 13(2), Article 39. <https://doi.org/10.3390/fi13020039>

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425-478.

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? *International Journal of Educational Technology in Higher Education*, 16, Article 39.