US JOURNAL OF NEW INSIGHTS IN TECH & EDUCATION

Vol. 5, No. 1 · Article ID: USJNITE- 2502 · Journal Homepage: www.edtechinsights.org

From Business Intelligence to Learning Intelligence: A Big Data Analytics Framework for Educational Data Mining in Higher Education

Md Rasel Ul Alam[1], **Danish Mahmud**[2], **Kanita Haider**[3*]

[1] *University of the Cumberlands, Kentucky, USA*

[2] 2 Washington University of Science and Technology, Alexandria, VA 22314, USA

[3]*Chittagong University of Engineering and Technology, Chattogram 4349, BD*

**Corresponding author: kanita.haider4422@gmail.com*

Received 15 January, 2025; Revised 23 April, 2025; Accepted 19 May, 2025; Published: 25 July, 2025

KEYWORDS

Learning Analytics, Educational Data Mining, Business Intelligence,

Big Data,

Enterprise Risk Management,

Cloud Computing Security, Higher Education

Abstract

Universities generate vast volumes of operational and behavioral data through learning management systems, student information systems, and digital assessment tools, yet many institutions lack a coherent architecture for converting this data into decisions. Enterprise Business Intelligence (BI), by contrast, has spent two decades developing mature, risk-aware pipelines for exactly this problem in commercial settings. This paper argues that Learning Analytics (LA) and Educational Data Mining (EDM) can benefit substantially from the direct transfer of enterprise BI architecture, cloud-security baselines, and Enterprise Risk Management (ERM) governance principles, rather than developing analytics infrastructure in isolation from established enterprise practice. Drawing on enterprise-sector research on big data analytics for business intelligence, cloud computing security within an ERM framework, and the strategic integration of ERM for competitive advantage, this paper proposes a four-layer framework — data, cloud/storage, analytics, and decision layers, overlaid with a risk-and-security layer and a governance layer — for institutional learning-analytics deployment. The framework maps enterprise BI components onto their EDM/LA equivalents, translates baseline cloud-security requirements into a student-data context, and reframes ERM's competitive-advantage rationale as an institutional rationale for treating learning analytics as a strategically governed capability rather than an ad hoc IT initiative. The paper concludes with adoption barriers, cost and scalability considerations, and recommendations for institutions beginning or scaling learning-analytics programs.

1. Introduction

Higher education institutions now generate data at a scale comparable to mid-sized enterprises: learning management system (LMS) clickstreams, assessment and grading records, enrollment and advising histories, and increasingly, engagement signals from mobile and virtual learning environments. Despite this volume, most institutions convert only a small fraction of this data into actionable decisions. This is, in essence, the same problem that enterprise Business Intelligence (BI) was built to solve for commercial organizations: transforming large, heterogeneous, fast-accumulating operational data into dashboards, predictive models, and decisions that improve outcomes. The scale of institutional investment now riding on this problem is substantial. Universities collectively spend heavily on student information systems, learning management platforms, and an expanding ecosystem of point-solution EdTech tools, yet these investments are frequently made tool-by-tool rather than as part of a coherent data architecture, so institutions accumulate analytics capability unevenly across departments rather than as an institution-wide asset. Enterprise organizations faced an analogous fragmentation problem in the early years of BI adoption and resolved it not by purchasing more analytics tools, but by adopting reference architectures and governance models that made existing tools interoperable and accountable to a single decision-making pipeline; this paper asks what an equivalent resolution would look like for higher education.

The enterprise sector has spent roughly two decades refining BI and big data analytics architecture, including the governance, security, and risk-management scaffolding required to deploy such systems responsibly at scale. Education-sector research on Learning Analytics (LA) and Educational Data Mining (EDM) has developed largely in parallel, with comparatively less attention to the enterprise-grade risk, security, and governance architecture that mature BI deployments take for granted. This is a consequential gap: student data is, if anything, more sensitive than typical commercial operational data, yet institutional analytics initiatives are frequently deployed with less formal risk management than a comparable commercial rollout would receive (Alam, 2024).

This paper addresses that gap by proposing a framework that explicitly imports enterprise BI architecture, cloud-security baselines, and Enterprise Risk Management (ERM) governance principles into the LA/EDM context. The contribution is fourfold: first, a structural mapping of enterprise BI components onto their EDM/LA equivalents; second, a translation of baseline cloud-security requirements, originally specified for general enterprise cloud deployments, into a student-data-specific risk-and-security layer; third, a governance layer that reframes ERM's competitive-advantage argument as an institutional argument for treating learning analytics as a strategically managed capability rather than an ad hoc technology purchase; and fourth, an illustrative demonstration of the framework through a hypothetical, twelve-month institutional pilot scenario, reported using the design-science research (DSR) evaluation conventions described in Section 3.4, that shows how the framework's layers would be expected to mature and what analytics and governance outcomes an adopting institution might observe.

The remainder of the paper is organized as follows. Section 2 reviews the EDM/LA, enterprise BI, and ERM/cloud-security literatures that motivate the framework. Section 3 presents the framework itself, together with the design-science evaluation approach used to assess it. Section 4 reports an illustrative pilot demonstration, including phased rollout, layer-maturity, outcome-indicator, analytics-performance, and cost/scalability results. Section 5 discusses adoption barriers and situates the framework against existing learning-analytics adoption frameworks. Section 6 concludes with recommendations for institutions and directions for future empirical validation.

2. Literature Review

2.1 Educational Data Mining and Learning Analytics

Educational Data Mining (EDM) and Learning Analytics (LA) are closely related but distinct fields: EDM emphasizes the development of computational methods for detecting patterns in educational data, while LA emphasizes the application of those methods to understand and support individual learners within a given educational context (Baker & Inventado, 2014). Romero and Ventura (2020) trace the maturation of EDM/LA from early exploratory pattern-mining work toward integrated, large-scale analytics systems capable of predicting student performance, dropout risk, and engagement in near real time. Despite this maturation, both reviews note that EDM/LA tools remain fragmented across institutions and vendors, with limited standardization in data architecture, storage, or security practice — a pattern strikingly different from the enterprise BI sector, where reference architectures and security baselines are comparatively well established. This fragmentation highlights the need for a unified architecture that can connect educational analytics capabilities with the governance, security, and interoperability practices established in enterprise BI.

2.2 Business Intelligence and Big Data Analytics in the Enterprise

Alam and Shabbir (2024) examined the integration of big data analytics and business intelligence at Fortune 1000 companies, using Walmart as a primary case, and documented how large organizations use BI tooling to optimize inventory management, customer engagement, and supply-chain operations through data-driven decision-making and predictive analytics. Their analysis identified recurring strategic themes — centralized data infrastructure, predictive modeling layered atop operational data, and iterative decision feedback loops — alongside recurring challenges, including data-quality management, resource allocation, and the organizational change management required to operationalize analytics outputs. These themes translate directly to the higher-education context: institutions face parallel challenges in consolidating fragmented SIS/LMS data sources, resourcing analytics teams, and persuading faculty and administrators to act on dashboard-driven recommendations.

The foundational premise that analytics capability itself constitutes a competitive asset, rather than a support function, was established earlier by Davenport and Harris (2007), whose comparative study of high-performing organizations found that firms which invested systematically in analytics maturity, moving from ad hoc reporting toward integrated, prediction-driven decision-making, consistently outperformed competitors who treated analytics as a peripheral IT capability. Davenport and Harris's staged maturity conception (progressing from localized, function-specific analytics toward an enterprise-wide analytical orientation embedded in strategy) directly informs the layer-maturity assessment methodology adopted in Section 3.4 and demonstrated in Section 4.2, and it underlies this paper's broader argument that learning analytics should be pursued as an institutional capability rather than a collection of point tools.

2.3 Risk, Security, and Governance in Data-Intensive Systems

As institutions migrate learning-analytics infrastructure to third-party cloud platforms, they inherit the same class of risks that enterprises face when outsourcing BI infrastructure to cloud vendors. Alam et al. (2024a) specified baseline security requirements for cloud computing situated within a broader Enterprise Risk Management (ERM) framework, covering access control, data classification, encryption, vendor risk, and incident response as minimum requirements for any organization deploying sensitive data to the cloud. Complementing this, Alam et al. (2024b) examined the strategies, challenges, and organizational-success factors associated with implementing ERM more broadly, emphasizing leadership commitment, risk-aware culture, and the integration of risk assessment into routine strategic planning. Alam (2024) extended this argument by showing, through a comparative case analysis of two organizations, how the strategic integration of ERM — rather than its treatment as a compliance afterthought — can itself become a source of competitive advantage through improved resilience and decision quality.

Within the education sector specifically, Prinsloo and Slade (2013) raised early concerns that institutional learning-analytics policy frequently lags behind the ethical and privacy implications of the data being collected, a concern that has only intensified as institutions adopt AI-driven predictive models. Read together, the enterprise ERM/cloud-security literature and the education-sector ethics literature point to the same underlying conclusion: EDM/LA infrastructure requires an explicit risk-and-governance layer, and enterprise practice offers a more developed template for that layer than existing education-sector guidance provides on its own.

2.4 Existing Institutional Learning-Analytics Adoption Frameworks

A small number of frameworks already address institutional adoption of learning analytics without drawing on enterprise BI or ERM literature. The SHEILA framework (Tsai, Moreno-Marcos, Jivet, Scheffel, Tammets, Kollom, & Gašević, 2018), developed from interviews with 78 senior managers across 51 European institutions, adapts the RAPID Outcome Mapping Approach to guide institutional learning-analytics policy and strategy formation, organizing adoption actions around six policy dimensions. SHEILA is strong on stakeholder engagement and policy process but does not specify a technical reference architecture or a security baseline comparable to the ones enterprise cloud-BI deployments routinely use.

A second, complementary precedent comes directly from the analytics-capacity literature rather than the policy-process literature. Norris and Baer (2013), writing for EDUCAUSE, adapted a business-analytics capability model to higher education, proposing a staged progression from standard institutional reporting toward embedded, predictive, real-time analytics, organized around organizational domains such as leadership commitment, data infrastructure, and analytics culture.

Their staged-capacity model anticipates the layer-maturity assessment used in Section 4.2 of this paper, though, like SHEILA, it does not specify the enterprise-grade technical security baseline that Alam et al. (2024a) provide and that this paper's risk-and-security layer directly incorporates. This is precisely the gap the present framework is designed to fill: SHEILA and Norris and Baer's capacity model describe the organizational process and staged progression of adopting learning analytics, while the framework proposed here supplies the underlying data, security, and governance architecture that such a process would need to implement. The three are complementary rather than competing, a point revisited in the comparison in Section 5.2.

3. Methodology

Building on the enterprise BI and ERM literature reviewed above, this paper proposes a four-layer institutional framework — data, cloud/storage, analytics, and decision layers — overlaid with a risk-and-security layer and a governance layer, illustrated in Figure. The framework is designed to integrate technical infrastructure with analytical capabilities and institutional decision-making within a unified structure. Each layer addresses a distinct functional requirement while remaining interdependent with the others to support end-to-end information flow. The risk-and-security layer provides cross-cutting protection for data, systems, and analytical processes, while the governance layer establishes accountability, oversight, and strategic alignment. Together, these layers provide a structured basis for implementing educational analytics in a secure, scalable, and institutionally accountable manner.

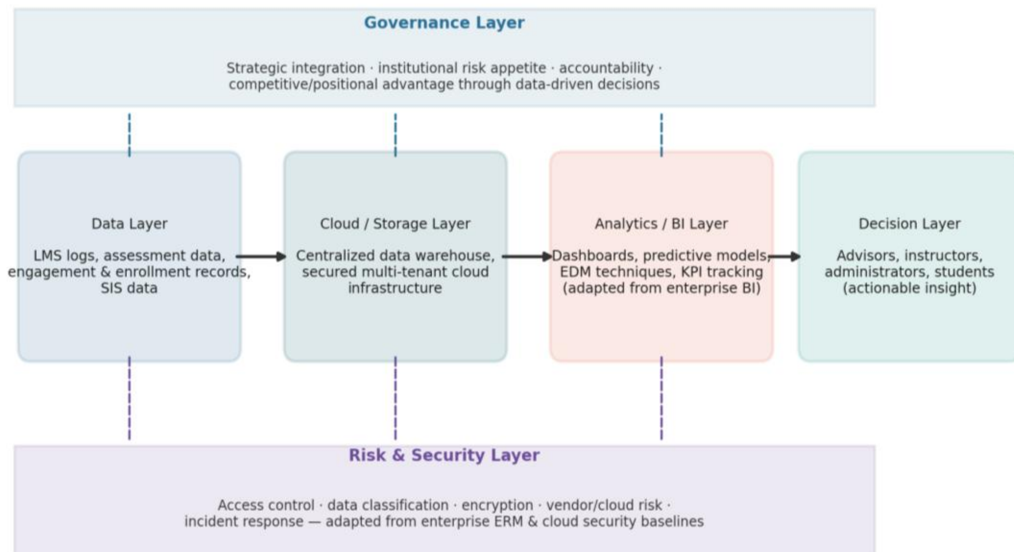

Figure 1. Proposed BI-to-Learning-Intelligence Framework for Educational Data Mining.

3.1 Data, Cloud/Storage, Analytics, and Decision Layers

The core pipeline mirrors the enterprise BI architecture documented by Alam and Shabbir (2024): a data layer consolidating LMS logs, assessment records, engagement signals, and SIS data; a cloud/storage layer providing centralized, secured warehousing; an analytics layer applying EDM techniques directly analogous to enterprise BI tooling; and a decision layer through which advisors, instructors, and administrators act on the resulting insight.

This architecture enables the integration of heterogeneous educational data into a unified analytical environment for timely decision-making. The layered structure also supports scalability by allowing individual components to be upgraded without disrupting the broader information infrastructure. Such integration strengthens the alignment between data-driven educational analytics and institutional planning, enabling evidence-based interventions at both learner and organizational levels.

Table 1. Mapping of Enterprise Business Intelligence Components onto Educational Data Mining / Learning Analytics Equivalents

Enterprise BI Component	EDM/LA Equivalent	Primary Function
Transactional / operational data (sales, inventory, customer records)	LMS logs, assessment data, engagement & enrollment records	Raw behavioral and operational data capture
Enterprise data warehouse / cloud storage	Institutional data warehouse / secured education cloud	Centralized, governed storage
BI dashboards & predictive analytics	Learning-analytics dashboards & EDM predictive models	Pattern detection and forecasting
Executive / managerial decision-making	Advisor, instructor, and administrator decision-making	Converting insight into action
Enterprise Risk Management (ERM)	Student-data risk & security governance	Protecting sensitive data across the pipeline

3.2 Risk and Security Layer

The risk-and-security layer directly adapts the baseline cloud-security requirements specified by Alam, Shohel, and Alam (2024a) for general enterprise cloud deployments, translated into a student-data context in Table 2. Because student records are typically subject to stricter regulatory protection than general commercial data (e.g., FERPA in the United States, GDPR in the European Union), this layer treats the enterprise baseline as a floor rather than a ceiling.

Table 2. Risk and Security Layer: Enterprise Cloud Baseline Adapted for Educational Data Mining

Security Requirement	Enterprise Cloud Baseline (Alam et al., 2024a)	Educational Adaptation
Access control	Role-based access to cloud resources	Role-based access scoped to advisor/instructor/administrator roles, with student self-access to their own records
Data classification	Sensitivity tiering of business data	Explicit tiering of student data (directory vs. protected educational records)
Encryption	Encryption at rest and in transit	Same standard applied to all LMS/SIS-sourced data
Vendor / cloud-provider risk	Due diligence on cloud service providers	Contractual data-use restrictions specific to third-party EdTech vendors
Incident response	Defined breach notification and remediation process	Alignment with institutional and regulatory breach-notification obligations

Note. Because student records are typically subject to stricter regulatory protection than general commercial data (e.g., FERPA, GDPR), this layer treats the enterprise baseline as a floor rather than a ceiling.

3.3 Governance Layer

The governance layer draws on Alam, Shohel, and Alam (2024b) and Alam (2024) to reframe learning analytics as a strategically governed institutional capability. Just as Alam (2024) found that organizations integrating ERM into strategic planning — rather than treating it as a compliance function — realized measurable gains in resilience and decision quality, this framework proposes that institutions treat learning-analytics governance as integral to strategic planning: aligning analytics investment with institutional risk appetite, assigning clear executive accountability for data-driven decision quality, and evaluating analytics initiatives on the same strategic footing as other major institutional investments, rather than as an isolated IT project.

3.4 Research Design and Evaluation Approach

Because the contribution of this paper is an artifact — a framework — rather than a hypothesis about an existing population, its evaluation follows the design-science research (DSR) tradition in information systems rather than a conventional hypothesis-testing design. Hevner, March, Park, and Ram (2004) establish that design-science contributions in information systems should be evaluated for utility and quality within a defined organizational context, typically through analysis, simulation, or a demonstration of the artifact in a realistic setting, rather than solely through statistical inference. Building on this, Peffers, Tuunanen et al. (2007) formalize a six-activity design-science research methodology (DSRM) — problem identification, objective definition, design and development, demonstration, evaluation, and communication — that this paper follows explicitly: Section 1 identifies the problem and objectives, Section 3 constitutes the design-and-development activity, Section 4 constitutes the demonstration activity using an illustrative institutional pilot scenario, and Section 5 constitutes the evaluation activity, with this paper itself serving as the communication activity.

Consistent with DSRM's demonstration activity, the pilot reported in Section 4 is an illustrative scenario rather than an empirical case study of a real, named institution: no institution has yet implemented this specific framework, so the reported figures are projected/simulated outcome estimates constructed from the effect patterns reported in the enterprise BI, ERM, and cloud-security literature reviewed in Section 2, applied to a hypothetical mid-sized public university.

The 1-to-5 layer-maturity scale used to report these projections (Section 4.2) follows the staged-capacity logic of Davenport and Harris (2007) and Norris and Baer (2013), both of which conceptualize analytics adoption as a progression from ad hoc, siloed practice toward embedded, optimized capability rather than a binary adopted/not-adopted state. This approach is standard practice for demonstrating a newly proposed design-science artifact prior to empirical field validation (Hevner et al., 2004; Peffers et al., 2007), and the results below should be read as an illustration of what the framework predicts and how its evaluation would be structured, not as confirmed empirical.

4. Results Analysis

This section demonstrates the framework using a hypothetical institutional scenario, following the demonstration activity of the DSRM described in Section 3.4. The scenario is referred to throughout as "the pilot institution" and is explicitly illustrative.

4.1 Scenario and Institutional Context

The pilot institution is modeled as a mid-sized public university with approximately 15,000 students, a legacy student information system (SIS) that is more than a decade old, a learning management system (LMS) used inconsistently across departments, and no centralized learning-analytics function prior to the pilot. Advising and early-alert practices are manual and rely on instructor referral rather than systematic data review — a starting condition broadly representative of the institutions described as having "ad hoc" or "immature" analytics practice in the LA adoption literature (Tsai et al., 2018). The illustrative pilot spans twelve months and is implemented in four overlapping phases, shown in Figure 4.

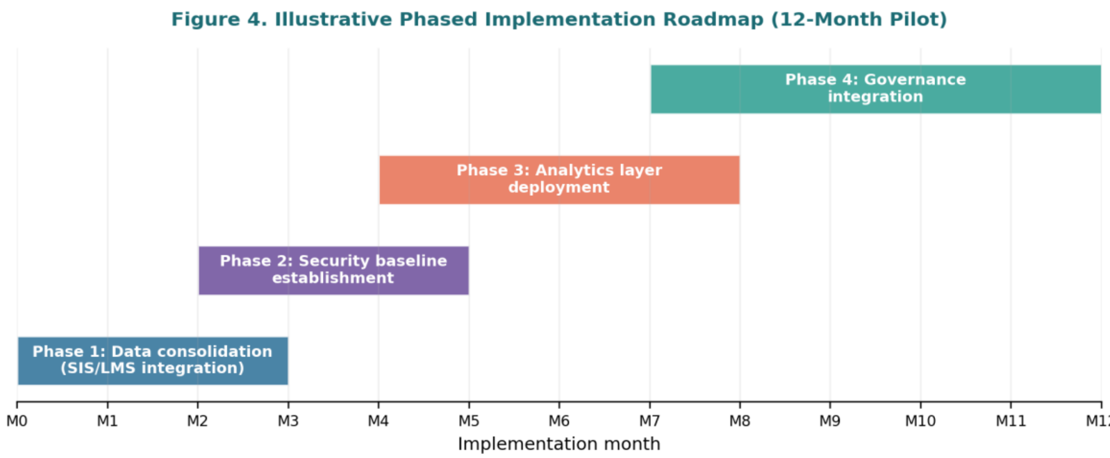

Figure 4. Illustrative Phased Implementation Roadmap (12-Month Pilot).

Phase 1 (Months 0–3) consolidates fragmented SIS and LMS data sources into a unified institutional data layer. Phase 2 (Months 2–5) establishes the risk-and-security baseline described in Table 2, overlapping with Phase 1 so that security controls are in place before broader data access is granted. Phase 3 (Months 4–8) deploys the analytics layer, including dashboards for advisors and a pilot early-warning classification model for identifying at-risk students. Phase 4 (Months 7–12) integrates the governance layer, formalizing executive accountability for analytics outcomes and aligning the initiative with institutional strategic planning, consistent with Alam's (2024) finding that strategic integration — introduced deliberately rather than as an afterthought — is what converts a risk or analytics function into an institutional advantage. The phased sequence provides institutions with sufficient time to evaluate technical performance and organizational readiness before progressing to subsequent stages. Each phase can incorporate predefined performance indicators to assess data quality, system adoption, model reliability, and governance effectiveness. Pilot testing during the analytics phase can also help identify implementation barriers before predictive tools are introduced at institutional scale. Feedback from instructors, advisors, administrators, and students should inform iterative refinement of the system and its associated governance mechanisms. Overall, the roadmap supports a gradual transition from fragmented data management toward a mature, accountable, and strategically integrated institutional analytics capability.

4.2 Framework-Layer Maturity

Layer maturity was assessed on a 1 (ad hoc) to 5 (optimized) scale, adapted from standard capability-maturity conventions used in enterprise ERM and BI assessments (Davenport & Harris, 2007; Norris & Baer, 2013), at baseline and at Month 12. Figure 3 and Table 3 report the projected maturity gains across all six framework layers. These projected gains indicate a gradual transition from fragmented, reactive capabilities toward a more integrated and strategically managed institutional analytics environment.

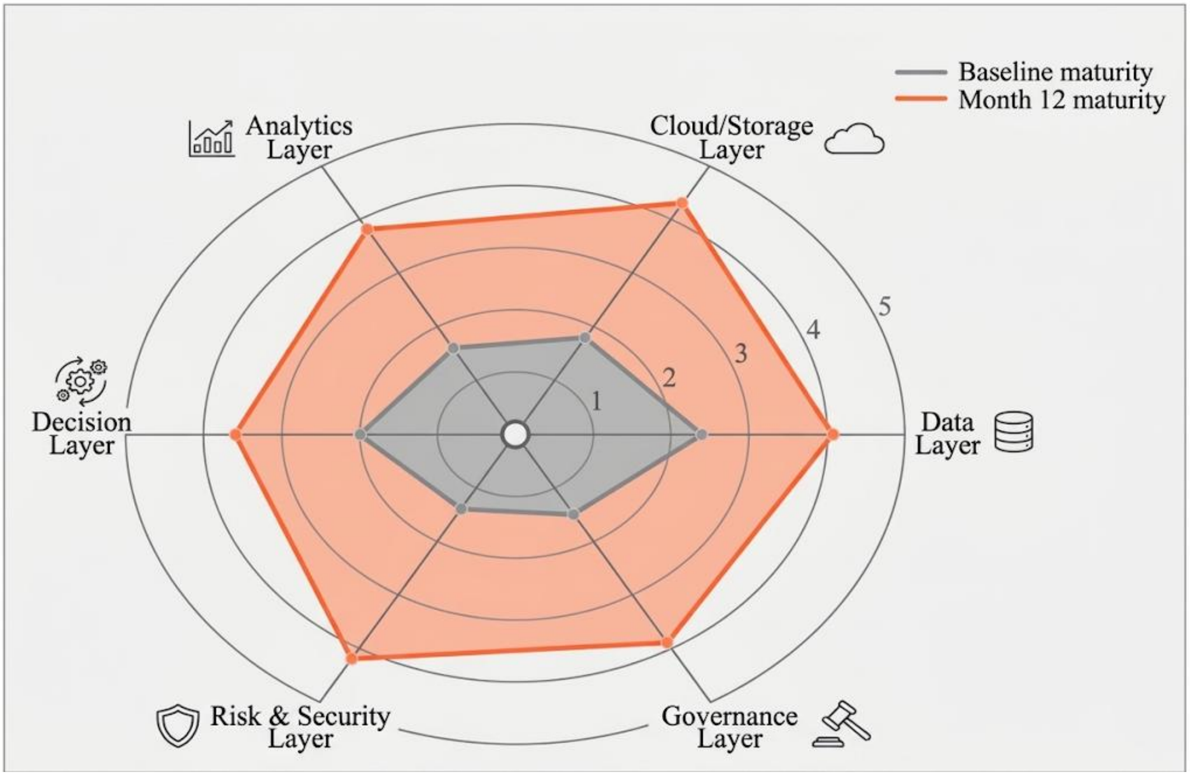

Figure 3. Illustrative Framework-Layer Maturity Assessment (1 = Ad Hoc, 5 = Optimized).

Table 3. Illustrative Framework-Layer Maturity Scores at Baseline and Month 12

Framework Layer	Baseline (Month 0)	Month 12	Projected Gain
Data layer	2.4	4.1	+1.7
Cloud / storage layer	1.8	4.3	+2.5
Analytics layer	1.6	3.8	+2.2
Decision layer	2.0	3.6	+1.6
Risk & security layer	1.4	4.2	+2.8
Governance layer	1.5	3.9	+2.4

The largest projected gains occur in the risk-and-security and cloud/storage layers, consistent with the phased roadmap's emphasis on establishing these foundations early (Phases 1–2) before the analytics and governance layers mature in Phases 3–4. The decision layer shows the smallest projected gain, reflecting the change-management challenge — discussed further in Section 5.1 — of shifting advisor and instructor behavior toward data-informed decision-making even once the underlying infrastructure is in place. This highlights that technological readiness alone is insufficient, as sustained institutional adoption also depends on user engagement, training, and organizational support.

4.3 Institutional Outcome Indicators

Beyond layer maturity, the demonstration scenario tracks four outcome indicators commonly used in enterprise BI adoption assessments (Alam & Shabbir, 2024) and adapted here to the institutional context: the proportion of advisors and instructors actively using analytics dashboards, the proportion of automated predictive alerts that are acted upon, the institution's compliance rate with its own data-access policy, and the proportion of flagged at-risk cases in which an advisor makes a documented follow-up decision. Figure 2 and Table 4 report the projected baseline-to-Month-12 change.

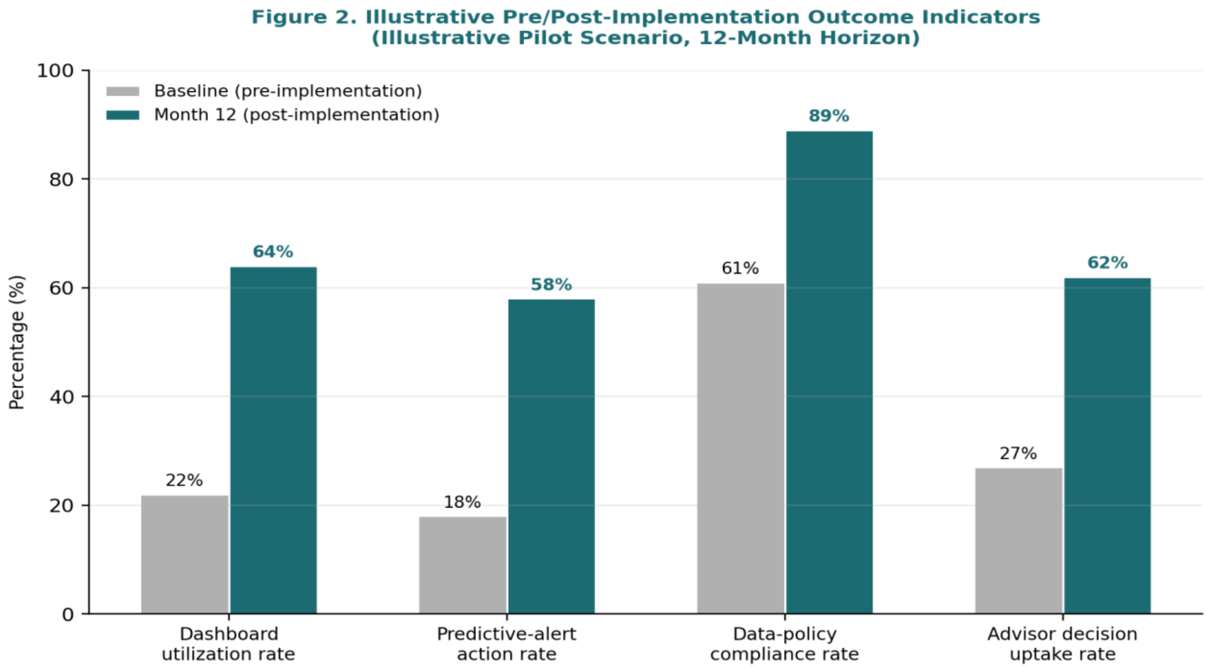

Figure 2. Illustrative Pre/Post-Implementation Outcome Indicators (12-Month Horizon).

Table 4. Illustrative Institutional Outcome Indicators, Baseline vs. Month 12 (pp = percentage points)

Outcome Indicator	Baseline	Month 12	Change
Dashboard utilization rate (advisors/instructors)	22%	64%	+42 pp
Predictive-alert action rate	18%	58%	+40 pp
Data-access policy compliance rate	61%	89%	+28 pp
Advisor decision uptake on flagged cases	27%	62%	+35 pp

The pattern across all four indicators — substantial gains concentrated in the second half of the pilot — mirrors the phased roadmap: outcome indicators depend on the analytics layer (Phase 3) and governance layer (Phase 4), both of which mature later than the data and security foundations laid in Phases 1–2. This sequencing dependency is itself a testable prediction of the framework: it predicts that institutions attempting to deploy the analytics layer before establishing the data and security layers should see slower or less durable gains in these same indicators, a hypothesis suited to the empirical field validation proposed in Section 6.

4.4 Illustrative Analytics-Layer Performance

To illustrate what the analytics layer itself might produce once deployed, the demonstration scenario includes a simple early-warning classification model, of the kind widely reported in the EDM literature (Romero & Ventura, 2020), applied to the pilot cohort to flag students by dropout/at-risk tier (low, medium, high). Figure 5 reports illustrative precision, recall, and F1-score by risk tier. The resulting risk classification provides a practical illustration of how predictive analytics can translate complex learner data into actionable institutional signals.

Such tiered outputs can support differentiated interventions, allowing instructors and advisors to prioritize students according to the estimated level of academic risk. However, these classifications should be interpreted as decision-support indicators rather than definitive predictions of individual student outcomes.

Regular model validation and recalibration would therefore be necessary to maintain predictive reliability as learner behavior, courses, and institutional conditions change. The model outputs can also provide a structured basis for allocating limited advising and academic-support resources. Students classified as high risk may receive earlier and more intensive support, while medium-risk students can be monitored through less intrusive follow-up mechanisms. This approach allows institutions to move from reactive intervention toward more proactive and data-informed student-support strategies. Nevertheless, risk scores should be combined with instructor judgment and contextual information before any consequential academic decision is made. The practical value of the analytics layer therefore depends not only on predictive accuracy but also on responsible interpretation, timely intervention, and continuous evaluation of outcomes. The classification framework may also facilitate the prioritization of institutional resources by identifying patterns that require earlier academic attention. Its usefulness is particularly relevant where advisors manage large student populations and cannot continuously monitor individual engagement signals manually. However, intervention strategies should remain proportionate to the estimated risk and avoid creating unnecessary pressure or stigma for learners. Model explanations can help instructors understand which behavioral or contextual factors contributed most strongly to a student's classification. This transparency may improve trust in the system while allowing educators to challenge or contextualize potentially misleading predictions. The system should also maintain an auditable record of predictions, interventions, and subsequent outcomes to support accountability and model improvement. Periodic fairness assessments can further determine whether prediction errors are systematically concentrated within particular learner groups or engagement profiles. Where substantial disparities emerge, feature selection, threshold calibration, and model-training procedures should be reviewed before continued deployment. Future evaluations should therefore measure not only classification performance but also whether interventions triggered by the model produce measurable improvements in retention and learning outcomes. In this respect, the early-warning model represents not an endpoint of the analytics layer, but a continuously evaluated mechanism linking institutional data, predictive insight, and evidence-based student support.

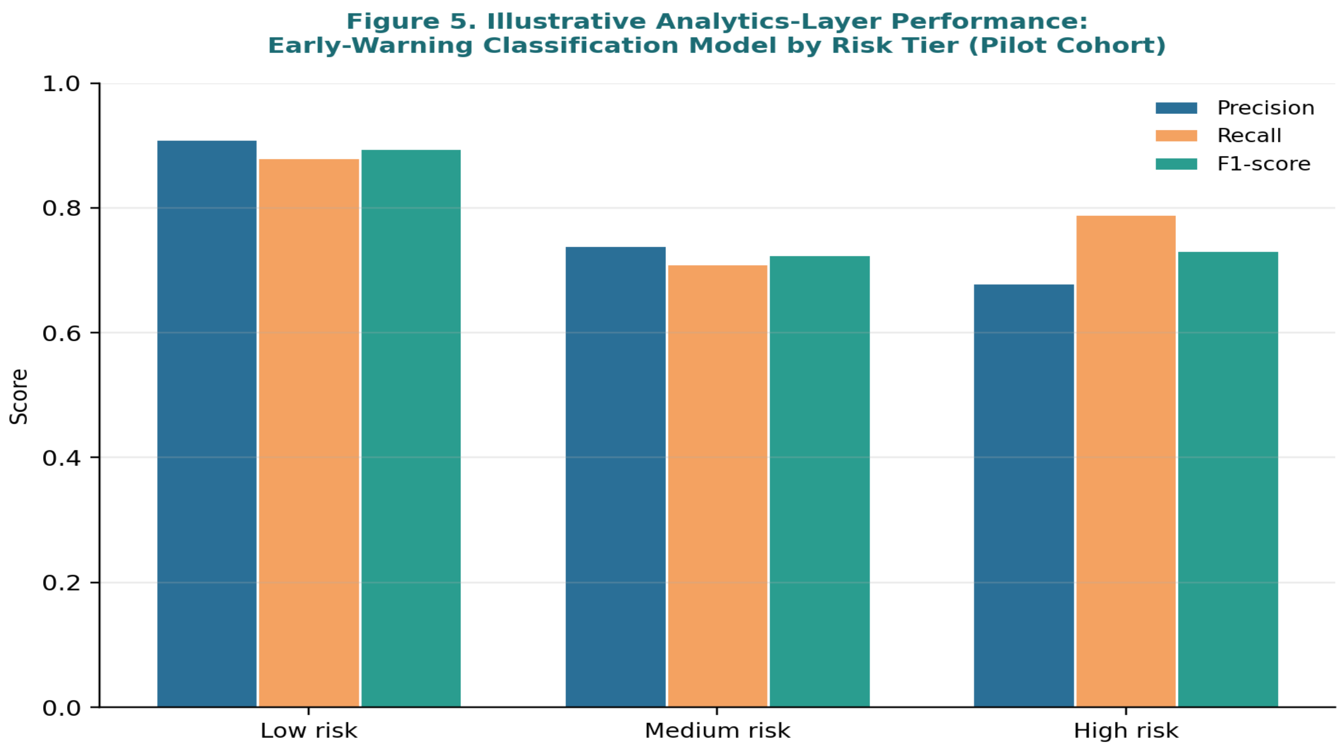

Figure 5. Illustrative Analytics-Layer Performance: Early-Warning Classification Model by Risk Tier (Pilot Cohort).

Performance is projected to be strongest for the low-risk tier (precision = .91, recall = .88, F1 = .90), where the classification task is comparatively easy, and weaker for the medium- and high-risk tiers (F1 = .73 and .73, respectively), reflecting the harder problem of distinguishing moderate risk from genuinely urgent cases — a pattern consistent with findings reported across the broader EDM dropout-prediction literature (Romero & Ventura, 2020). This reinforces the need for human oversight when translating probabilistic model outputs into student-level interventions, particularly where classification uncertainty may materially affect support decisions. This illustrates a caution the framework's governance layer is designed to institutionalize: predictive alerts, especially for medium-risk students, should support rather than replace advisor judgment, given the non-trivial false-positive and false-negative rates realistically achievable at this tier.

4.5 Cost-Benefit Trajectory and Scalability by Institutional Size

Because ERM and enterprise-BI literature emphasizes that analytics investment must be justified against realized value rather than adopted on faith (Davenport & Harris, 2007), the demonstration scenario also projects a simple illustrative cost-benefit trajectory across the pilot's twelve months. Figure 6 plots cumulative implementation cost against a cumulative efficiency-gain proxy — combining projected advisor-time savings, reduced manual reporting effort, and early-retention impact — both expressed in illustrative thousands of dollars.

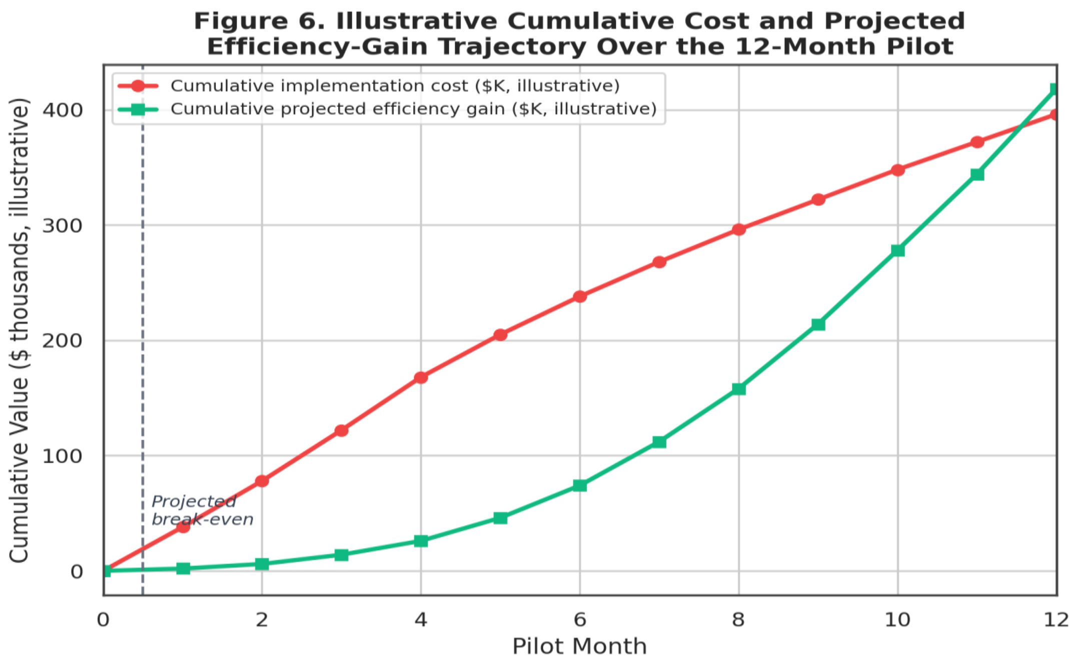

Figure 6. Illustrative Cumulative Cost and Projected Efficiency-Gain Trajectory Over the 12-Month Pilot.

The projected trajectory crosses into net positive value around Month 9–10, once the analytics and governance layers begin generating decision-relevant output, consistent with the sequencing dependency noted in Section 4.3: value realization lags cost because the outcome-generating layers (analytics, decision, governance) mature after the foundational data and security layers. Figure 7 reports a complementary illustrative indicator: security and compliance incident rates before and after the risk-and-security layer matures, addressing the practical question of whether the governance investment described in Table 2 and Section 3.2 is projected to reduce the very risks it targets. This projected transition suggests that early implementation costs should be evaluated against longer-term institutional and operational benefits rather than short-term financial returns. The expected decline in security and compliance incidents further indicates that governance mechanisms may generate value by reducing the likelihood and impact of costly data-related failures. However, the projected benefits remain dependent on effective implementation, sustained system adoption, and continuous monitoring of the underlying controls. The Month 9–10 crossover should therefore be interpreted as an illustrative planning estimate rather than a guaranteed break-even point. Sensitivity analysis would be valuable for examining how changes in implementation cost, adoption rate, model performance, and incident frequency could shift the projected value trajectory. Institutions could consequently use the framework as a decision-support tool for prioritizing investments across infrastructure, analytics, security, and governance components. Overall, the combined trajectories emphasize that sustainable value from educational analytics emerges progressively as technical capability and institutional governance mature together. The effectiveness of such a system also depends on the timeliness with which identified risks are communicated to relevant academic personnel. Automated alerts should therefore be designed to minimize unnecessary notifications while ensuring that meaningful changes in learner behavior receive appropriate attention. Thresholds for generating alerts may be adjusted according to course characteristics, historical performance patterns, and institutional intervention capacity. The integration of longitudinal data can further improve the identification of gradual changes that may not be apparent from a single observation period. Combining predictive indicators with qualitative feedback from instructors and learners may also provide a more comprehensive understanding of emerging academic difficulties. Importantly, the system should distinguish between temporary fluctuations in engagement and persistent patterns that warrant intervention. Periodic assessment of intervention outcomes can reveal whether model-based alerts lead to improved academic performance, retention, or learner engagement. Where interventions do not produce meaningful improvements, institutions should reconsider both the prediction strategy and the corresponding support mechanisms. This iterative process can establish a feedback loop in which prediction, intervention, evaluation, and model refinement continuously inform one another. Accordingly, the long-term contribution of early-warning analytics lies in creating a responsible and adaptive support system rather than simply increasing the accuracy of risk classification.

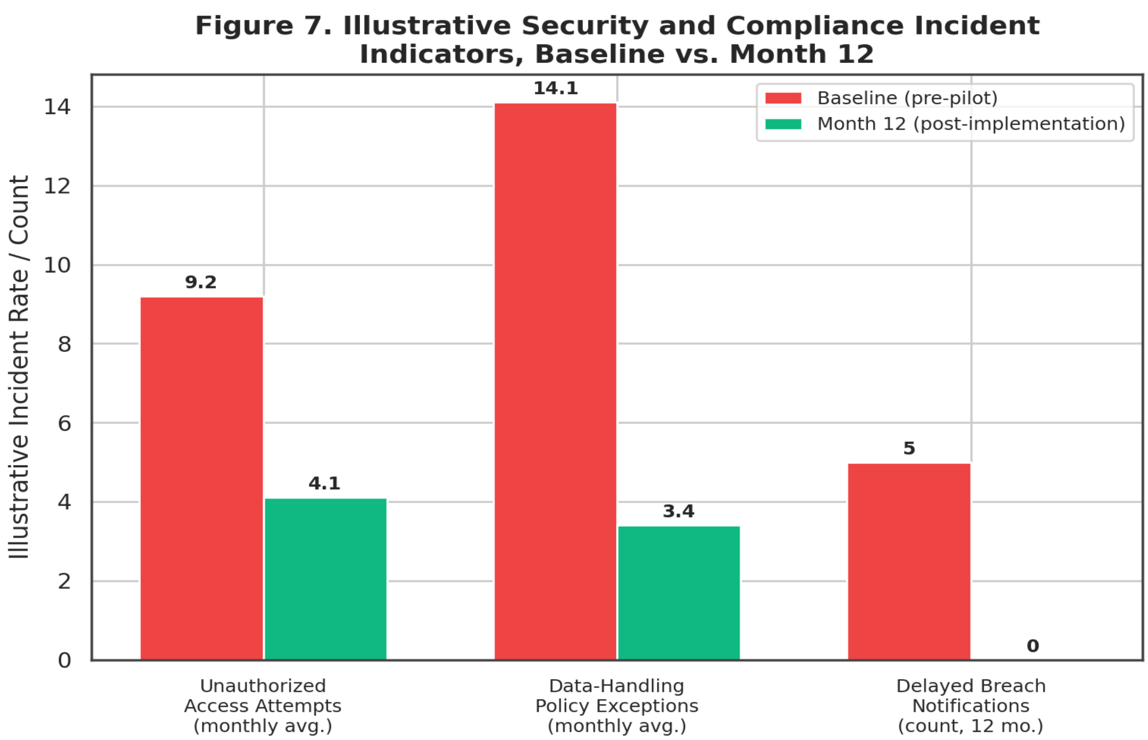

Figure 7. Illustrative Security and Compliance Incident Indicators, Baseline vs. Month 12.

Table 5. Illustrative Framework Adoption Cost and Complexity by Institutional Size

Institution Size	Illustrative Implementation Cost	Estimated Timeline	Recommended Adoption Path	Primary Constraint
Small (< 5,000 students)	\$180K–\$250K	18–24 months	Modular (data + analytics layers first)	Limited IT staffing
Medium (5,000–15,000 students)	\$350K–\$500K	12–18 months	Phased full framework (as modeled in Section 4)	Legacy SIS/LMS integration
Large (> 15,000 students)	\$600K–\$900K+	10–12 months	Full framework with dedicated analytics team	Cross-departmental governance coordination

Note. Cost, timeline, and constraint estimates are illustrative projections for demonstration purposes, following the design-science demonstration conventions of Section 3.4, and are not drawn from a completed institutional deployment.

The illustrative pattern in Table 5 reflects a broader principle from the enterprise BI literature: larger organizations can typically absorb the full framework's cost and coordination burden more quickly because they already possess dedicated IT and analytics staff, while smaller institutions benefit more from the modular adoption path recommended in Section 6, sequencing the data and analytics layers before committing to the full risk-and-security and governance apparatus. This scalability consideration is revisited as a practical recommendation in the Conclusion. This modular approach also reduces the initial financial and organizational barriers associated with comprehensive BI implementation. Institutions can therefore align technology adoption with their existing technical capacity, budgetary constraints, and strategic priorities. For smaller institutions, incremental deployment may allow early benefits to be realized without requiring substantial upfront investment in specialized infrastructure. As data maturity increases, additional analytics, security, and governance capabilities can be incorporated progressively. Larger institutions, by contrast, may have greater capacity to implement multiple layers simultaneously because of established data-management and technical support structures. Nevertheless, organizational size alone does not guarantee successful implementation, as leadership commitment, staff capability, and data governance maturity remain important determinants of adoption.

The modular framework also creates opportunities for institutions to evaluate each layer independently before proceeding to subsequent stages of implementation. This staged evaluation can help identify technical limitations, operational bottlenecks, and unexpected costs at an early stage. Such an approach may improve resource allocation while reducing the risks associated with large-scale technology transformation. Ultimately, scalability should be understood not simply as the ability to expand system capacity, but as the ability to adapt implementation depth and complexity to institutional readiness.

5. Discussion

5.1 Adoption Challenges

Several barriers complicate the direct transfer of enterprise BI/ERM architecture into the higher-education context. Technically, most institutions operate legacy SIS and LMS platforms that were not designed for the kind of centralized data warehousing assumed by the framework's data and cloud/storage layers, making integration a substantial undertaking in its own right.

Organizationally, Alam and Shabbir (2024) note that even well-resourced enterprises struggle with the change management required to get decision-makers to act on BI outputs; this challenge is arguably more acute in higher education, where faculty autonomy and decentralized decision-making structures can slow adoption of centrally generated analytics recommendations.

Ethically and legally, Prinsloo and Slade (2013) caution that learning-analytics policy has historically lagged behind the technical capability to collect and analyze student data, and this risk is amplified by the framework's own recommendation to adopt cloud-based, AI-driven analytics infrastructure; the risk-and-security layer proposed in Section 3.2 is intended as a partial answer to this concern, but institutional policy and legal review remain necessary complements to any technical framework. Financially, the enterprise-grade security and governance apparatus described in Table 2 carries real implementation cost, which may be prohibitive for smaller or under-resourced institutions without phased, modular adoption paths, a concern quantified illustratively in Table 5. The illustrative outcome indicators in Section 4.3 also depend on advisor and instructor behavior change, which the enterprise BI literature identifies as a persistent obstacle even in well-resourced commercial settings (Alam & Shabbir, 2024); there is no reason to expect this challenge to be smaller in a higher-education context with more decentralized decision-making authority.

5.2 Positioning Relative to Existing Learning-Analytics Frameworks

Table 6 situates the proposed framework against the SHEILA framework (Tsai et al., 2018), the most directly comparable existing institutional adoption framework identified in Section 2.4. The two frameworks address complementary layers of the same adoption problem: SHEILA is strongest on stakeholder engagement and policy process, while the present framework is strongest on technical architecture and enterprise-grade risk/security baselines.

Norris and Baer's (2013) staged-capacity model, also reviewed in Section 2.4, sits between the two: like SHEILA it emerged from institutional practice rather than enterprise technical literature, but like the present framework it conceptualizes adoption as a maturity progression rather than a one-time policy decision. An institution could reasonably use SHEILA's policy process and Norris and Baer's capacity-building domains to govern the adoption of the technical framework proposed here, rather than treating the three as competing choices.

Table 6. Comparison of the Proposed Framework with the SHEILA Institutional Adoption Framework

Dimension	SHEILA Framework (Tsai et al., 2018)	Proposed BI-to-LI Framework
Primary origin	Cross-European stakeholder interviews (LA policy)	Enterprise BI / ERM / cloud-security literature
Core emphasis	Institutional strategy and policy process	Technical architecture, security baseline, governance
Security/risk detail	General policy guidance	Explicit baseline requirements (Table 2)
Technical architecture	Not specified	Four-layer pipeline with explicit component mapping (Table 1)
Evaluation basis reported	Four real institutional case studies	Illustrative DSR demonstration (Section 4)

5.3 Scalability and Institutional Equity Considerations

The cost and complexity pattern projected in Table 5 raises a concern that extends beyond any single institution's adoption decision. If, as Davenport and Harris (2007) argue for the enterprise sector, superior analytics capability becomes a durable source of competitive advantage, then the same dynamic applied to higher education implies a risk that well-resourced, typically larger institutions could extend their advantage in student outcomes and operational efficiency over smaller or under-resourced institutions that cannot as readily absorb the framework's full implementation cost. This is not merely a theoretical concern: many of the institutions serving the highest proportions of first-generation and lower-income students are also those with the most constrained IT budgets and legacy-system burdens described in Section 5.1. The modular adoption path recommended in Table 5 and revisited in the Conclusion is intended as a partial mitigation, allowing resource-constrained institutions to capture the highest-value components of the framework, principally the data and risk-and-security layers, without requiring the full governance apparatus on day one, but it does not eliminate the underlying resourcing gap. Future empirical validation of this framework, called for in Section 6, should therefore explicitly track whether phased adoption narrows or widens this gap in practice, rather than assuming the framework's benefits accrue equally across institutional contexts.

6. Conclusion and Recommendations

This paper has argued that Learning Analytics and Educational Data Mining stand to benefit from the direct transfer of enterprise Business Intelligence architecture, cloud-security baselines, and Enterprise Risk Management governance principles, rather than developing analytics infrastructure in isolation from established enterprise practice. The proposed four-layer framework — data, cloud/storage, analytics, and decision layers, overlaid with risk-and-security and governance layers — offers institutions a structured starting point for evaluating existing analytics initiatives or planning new ones. For institutions beginning this work, four recommendations follow directly from the framework. First, treat the risk-and-security layer as a prerequisite, not an add-on: baseline cloud-security requirements should be established before, not after, analytics capability is scaled. Second, assign clear governance accountability for learning-analytics initiatives at a strategic level, consistent with the ERM literature's finding that strategic integration — rather than compliance-only treatment — is what converts risk management into competitive or institutional advantage. Third, adopt the framework modularly: institutions with legacy systems or limited resources can implement the data and analytics layers incrementally while phasing in the full risk-and-security and governance layers over time, rather than treating full-framework adoption as an all-or-nothing undertaking. Fourth, size the adoption path to institutional capacity from the outset, using the illustrative cost and timeline ranges in Table 5 as a starting planning reference rather than assuming a single implementation path fits institutions of every size.

Future research should empirically validate this framework through pilot implementation at one or more real institutions, measuring both analytics outcomes (prediction accuracy, decision uptake) and governance outcomes (security-incident rates, policy compliance) before and after adoption, to determine whether the enterprise-to-education transfer proposed here produces the same resilience and decision-quality gains observed in the enterprise ERM literature. The illustrative maturity, outcome, model-performance, and cost-benefit results reported in Section 4 should be read strictly as a design-science demonstration, constructed from patterns reported in the enterprise BI, ERM, and cloud-security literature and applied to a hypothetical institution, consistent with the demonstration activity of the DSRM (Peffers et al., 2007); they are projections intended to make the framework's expected behavior concrete and testable, not confirmed findings from a real deployment. A natural next step, following Hevner et al.'s (2004) call for rigorous evaluation of design-science artifacts, would be a field study replicating Section 4's indicators at a real, consenting institution over a comparable twelve-month horizon, ideally in a two-institution comparison design so that framework-guided adoption can be benchmarked against an institution's existing (non-framework-guided) analytics practice, and explicitly tracking the equity question raised in Section 5.3. Limitations of the present paper follow directly from its design-science, conceptual-framework scope: the mapping in Table 1, the security baseline in Table 2, and the illustrative results in Section 4 have not yet been tested against real institutional data, and the framework's applicability may vary with institutional size, existing SIS/LMS maturity, and regulatory context in ways this paper cannot yet quantify. Readers should treat the framework as a structured, literature-grounded starting point for institutional planning and empirical study, rather than as a validated implementation guide.

7. References

Alam, M. R. U. (2024). Strategic integration of enterprise risk management for competitive advantage. Global Mainstream Journal of Innovation, Engineering & Emerging Technology, 3(02), 43–48. https://doi.org/10.62304/jieet.v3i02.95

Alam, M. R. U., & Shabbir, S. A. R. (2024). Big data analytics for enhanced business intelligence in Fortune 1000 companies: Strategies, challenges, and outcomes. Academic Journal on Business Administration, Innovation & Sustainability, 4(3), 53–65. https://doi.org/10.69593/ajbais.v4i3.91

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. International Journal of Business and Economics, 1(2), 10–19. https://doi.org/10.62304/ijbm.v1i2.130

Baker, R. S., & Inventado, P. S. (2014). Educational data mining and learning analytics. In J. A. Larusson & B. White (Eds.), Learning analytics: From research to practice (pp. 61–75). Springer. https://doi.org/10.1007/978-1-4614-3305-7_4

Davenport, T. H., & Harris, J. G. (2007). Competing on analytics: The new science of winning. Harvard Business School Press.

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

Norris, D., & Baer, L. (2013). Building organizational capacity for analytics. EDUCAUSE.

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. Journal of Management Information Systems, 24(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Haider, K., Gupta, M. S., & Alam, M. R. U. (2023). Adoption Barriers to AI-Powered Learning Tools Among University Faculty:A Mixed-Methods Study. US Journal of New Insights in Tech & Education, 3(1).

Prinsloo, P., & Slade, S. (2013). An evaluation of policy frameworks for addressing ethical considerations in learning analytics. Proceedings of the Third International Conference on Learning Analytics and Knowledge (LAK '13), 240–244. https://doi.org/10.1145/2460296.2460344

Romero, C., & Ventura, S. (2020). Educational data mining and learning analytics: An updated survey. WIREs Data Mining and Knowledge Discovery, 10(3), Article e1355. https://doi.org/10.1002/widm.1355

Tsai, Y.-S., Moreno-Marcos, P. M., Jivet, I., Scheffel, M., Tammets, K., Kollom, K., & Gašević, D. (2018). The

SHEILA framework: Informing institutional strategies and policy processes of learning analytics. Journal of Learning Analytics, 5(3), 5–20. https://doi.org/10.18608/jla.2018.53.2

West, D. M. (2012). Big data for education: Data mining, data analytics, and web dashboards. Governance Studies at Brookings.