International Journal of Science and Research Archive
eISSN: 2582-8185
Cross Ref DOI: 10.30574/ijsra
Journal homepage: https://ijsra.net/

(CASE REPORT)

Check for updates

BALANCING ACCESS AND PRIVACY: REGULATORY GAPS, AUTHENTICATION, AND DATA SECURITY IN DIGITAL PRESERVATION OF INSTITUTIONAL RECORDS

Afua Asantewaa Asante [1,*], Kwaku Gyamfi Boamah [1] and Ranjith Kumar Patil [2]

[1] College of Computing – Cybersecurity, Grand Valley State University, USA.
[2] College of Computing - Applied Computer Science, Grand Valley State University, USA.
* Corresponding Author

ORCID Details
Afua Asantewaa Asante: https://orcid.org/0009-0001-5634-5759
Kwaku Gyamfi Boamah: https://orcid.org/0009-0002-8385-3684

International Journal of Science and Research Archive, 2026, 20(02), 311–322

Publication history: Received on 04 July 2026; revised on 12 August 2026; accepted on 14 August 2026

Article DOI: https://doi.org/10.30574/ijsra.2026.20.2.1558

ABSTRACT

Digital preservation is often framed as a technical archival problem, yet the governance choices surrounding it carry direct consequences for information security and data privacy. This study draws on qualitative case-study data from a Ghanaian educational institution to examine how the absence of enforced legal and regulatory frameworks for digitized records shapes user authentication practices, data security, and privacy outcomes in a live school management information system. Twenty-five stakeholders, including parents, teachers, facilitators, and administrators, were interviewed about their experiences with user access, authentication, and data handling. Findings show that authentication mechanisms, including unique login credentials, QR-code verification, and digital signatures, were valued by users primarily as privacy safeguards, and that stakeholders explicitly requested encryption and non-disclosure of data to third parties, despite the absence of a specific institutional data-protection policy referencing Ghana's regulatory framework. The findings are discussed in relation to Ghana's Data Protection Act, 2012, and the Public Records and Archives Administration Act, 1997, alongside the growing role of artificial-intelligence-assisted compliance monitoring and cloud security posture management. The study concludes that in resource-constrained institutional settings, privacy protection is being improvised by end users and administrators in the absence of formal governance, a pattern with implications for institutions well beyond the case examined here.

Keywords: Digital Preservation; Data Privacy; Information Security; User Authentication; Data Governance

1 INTRODUCTION

Institutions that digitize their records face a governance question that is easy to overlook amid the operational benefits of digitization: who is allowed to see what, on what authority, and under what safeguards. Digital preservation literature has traditionally emphasized the technical durability of digital records - format obsolescence, media fragility, migration strategies - while treating access control and privacy as secondary concerns [1]. Yet as records move from filing cabinets

into cloud-connected information systems, the question of who can access a child's academic record, a parent's contact details, or a staff member's login credentials becomes a live security and privacy issue, not merely an archival one.

This paper revisits data collected from a broader thesis-length study of digital preservation and data management at Divine International College (DIC), Accra, and isolates the governance dimension for focused treatment: the extent to which the non-enforcement of legal and regulatory frameworks for digitized records in Ghana shapes authentication practice, data security, and user privacy expectations in a live information system. Where the original study treated this as one of three research objectives alongside information literacy and infrastructure cost, this paper treats it as a self-contained governance and security case study, in dialogue with contemporary data-protection, cybersecurity, and cloud-security literature.

This paper makes the following contributions. First, it provides one of the few empirically grounded accounts of how authentication and privacy expectations are actually experienced and articulated by non-specialist stakeholders - parents, teachers, procurement staff - in a live, resource-constrained institutional information system, rather than by IT professionals or compliance specialists. Second, it demonstrates a concrete, replicable method for mapping such informally-expressed governance instincts onto formal concepts (identity and access management, encryption policy, third-party data-sharing agreements, access review) via the comparative table presented in Section 3.3, offering other researchers a template for similar cross-institutional or cross-national comparisons. Third, it connects this small-institution case, for the first time, to the recent enterprise- and healthcare-scale cloud security literature on posture management, workload protection, and confidential computing, arguing that the underlying governance problem is scale-invariant even where the available tooling is not. Fourth, it proposes a minimum viable governance checklist (Section 3.6) and a sequenced AI-assisted research agenda (Section 3.4) explicitly designed for institutions that cannot adopt enterprise-grade frameworks wholesale, addressing a gap between what the security literature typically recommends and what resource-constrained institutions can plausibly sustain.

The remainder of this paper is organized as follows. Section 3.1 reviews literature on digital preservation governance, authentication as a privacy mechanism, and emerging AI- and cloud-based governance tooling. Section 2 describes the qualitative methodology, including the research paradigm, sampling strategy, and validity procedures. Section 3.2 presents findings from stakeholder interviews. Section 3.3 discusses the implications of these findings for governance theory and practice. Section 3.4 proposes directions for AI-assisted compliance research. Section 3.5 offers practical recommendations. Section 3.6 discusses limitations, and Section 3.7 concludes.

2 MATERIAL AND METHODS

2.1 Research paradigm

This paper draws on qualitative case-study data originally gathered using an interpretivist and constructivist research paradigm. Interpretivism holds that human behavior is too complex to be studied through probabilistic, positivist models, and that knowledge is instead created by interpreting the meanings people attach to behaviors and events; results under this paradigm are contextually valid rather than broadly generalizable. Constructivism, similarly, holds that individuals construct their own understanding of reality through experience and reflection, and seeks the "why" of events rather than the "how." Both paradigms were judged appropriate for a study of digital preservation's impact on data management because the phenomenon under study - how institutional actors experience and respond to authentication, privacy, and security demands - is inherently a matter of subjective interpretation rather than an object that can be measured in isolation from the people experiencing it.

2.2 Case and sampling strategy

Divine International College, an educational institution in Accra, Ghana, that uses the PROCARE school management information system, was selected as the case using a maximal-variation, purposive and convenient sampling strategy. Rather than a large probabilistic sample, the study prioritized depth: a smaller number of information-rich cases was judged to yield a better understanding of the phenomenon than a larger, more superficial sample would, consistent with established guidance that the ability to provide an in-depth picture diminishes as more sites or individuals are added to a qualitative case study. The sampling frame comprised facilitators, parents, management members, procurement officers, teachers, IT associates, and administrators, for a target sample size of 27. Twenty-five of these 27 targeted stakeholders (83.3%) were ultimately interviewed; two DIC personnel could not be reached despite repeated scheduling attempts.

2.3 Data collection and analysis

Data were collected through structured interviews and analyzed using a first-order/second-order/aggregate-dimension coding approach, producing two theoretical dimensions directly relevant to this paper: "user access" (comprising user agreement, identification, authorization, privacy, and safety) and "robust system" (comprising data accuracy, data security, and data privacy). This paper re-examines that coded data specifically through a governance, cybersecurity, and privacy-law lens, rather than the general digital-preservation lens of the original study.

2.4 Reliability and validity procedures

Several procedures were used to support the reliability and validity of the underlying data. A research supervisor reviewed the interview guide for clarity prior to fieldwork. Respondents received advance notice of interviews to allow time to prepare. Interview data were coded and inspected systematically to reduce transcription errors and omissions. To support construct validity, a pilot round of questions was administered prior to the main data collection to check that constructs were clearly defined and that dimensions related to one another as expected. To support internal validity, interviews were recorded, transcribed, and returned to interviewees for member-checking prior to analysis, reducing the risk that findings reflect only the researchers' own interpretation. External validity - the extent to which findings can be expected to generalize - is addressed further in the Limitations section below, since a single-institution qualitative case study of this kind carries specific, known constraints on generalizability that should not be minimized.

2.5 Reflexivity and positionality

Qualitative case-study research of this kind is shaped by the researchers' own position relative to the setting studied, and this should be stated rather than left implicit. The parent study was conducted by researchers with an existing academic and professional interest in information systems and records management, interviewing stakeholders at an institution the research team had established access to. This relationship likely eased access and candor - stakeholders spoke openly about difficulties and failures alongside successes - but it also means the researchers were not fully external, disinterested observers, and the interpretive coding performed on the interview data reflects the research team's own analytic categories ("user access," "robust system") rather than categories the stakeholders themselves would necessarily have generated independently. Member-checking, described above, was used specifically to guard against the risk that these researcher-imposed categories distorted what stakeholders actually said, by returning transcripts to interviewees for confirmation before analysis proceeded.

2.6 Interview protocol themes

The interview protocol underlying this dataset was organized around open-ended questions addressing, in sequence: the respondent's role and relationship to the PROCARE system; their experience of gaining and using access to the system; their understanding of what happens to the data they or their dependents provide; their views on the security and privacy of that data; and any concerns or suggestions for improvement. Questions were deliberately open-ended rather than checklist-style, on the theory that a checklist format ("Does the system use encryption? Yes/No") would have elicited compliance-oriented answers rather than the candid, unprompted governance vocabulary - "it should be encrypted," "not shared to any third party" - that turned out to be the most analytically useful material in the dataset. This design choice explains why respondents raised specific technical expectations unprompted rather than simply confirming or denying a list of controls presented to them.

3 RESULTS

3.1 Literature Context

3.1.1 Digital preservation and the governance gap

Digital preservation refers to the strategies organizations use to maintain the reliability, authenticity, and long-term accessibility of digital information [1]. A substantial body of work has documented that the legal and policy infrastructure meant to govern this activity lags behind the technology itself. Tsvuura and Ngulube [2], studying digitization at two Zimbabwean universities, found that record-keeping legislation had not kept pace with technological change. Asogwa [3] and Mukred et al. [4] similarly note that most Sub-Saharan African countries lack dedicated legislation for electronic records and archives, leaving institutions to manage sensitive digitized information without a clear compliance backstop.

Ghana is illustrative rather than exceptional. The Public Records and Archival Administration Act, 1997 (Act 535) obliges public offices to maintain proper record-keeping and disposal practices, but the thesis this paper draws on found no institutional evidence that this legislation was being actively enforced or referenced in day-to-day digitization decisions at DIC. Ghana also has a dedicated privacy statute - the Data Protection Act, 2012 (Act 843) -administered by

the Data Protection Commission, which sets out principles for lawful processing, consent, and data-subject rights. Neither piece of legislation surfaced spontaneously in stakeholder interviews; instead, privacy and security were discussed by respondents in operational, common-sense terms - "not shared to any third party," "encrypted," "unique login" - that map onto legal obligations without being explicitly derived from them.

3.1.2 Authentication as a privacy-protective mechanism

A growing security literature treats authentication not merely as an access-control gate but as a direct privacy safeguard. Authentication research argues that authentication protects privacy precisely because it guarantees that only authorized users reach sensitive information or system functions; by validating identity, authentication reduces the risk of data breaches and unauthorized disclosure [5]. Related work on authentication measures shows that stronger mechanisms- multi-factor authentication and biometrics in particular - reduce the likelihood of illegitimate access to sensitive data, which in turn keeps personal data unavailable to unauthorized parties and protects user privacy as a downstream effect [6].

This reframes a common assumption in digital preservation practice: that security and privacy are separable concerns, with security addressing system integrity and privacy addressing data subject rights. The authentication literature instead treats them as mechanically linked - the same control (verifying identity before granting access) serves both goals simultaneously. This paper's findings support that linkage empirically.

3.1.3 Data accuracy, system design, and security

Security outcomes are also shaped by data quality and system architecture. Dika and Nowostawski [7] found that source-code and configuration vulnerabilities in digital systems put user data at risk in ways closely tied to how accurately and consistently the underlying system is built and maintained; accurate, well-structured data is less prone to the inconsistencies that create security vulnerabilities. Robust system design is understood in the wider security literature to incorporate firewalls, intrusion detection, and regular security audits to guard against unauthorized access and cyberattack [8,9]. Read together, this literature suggests that data governance in digitized record systems operates on at least three interacting layers: identity verification (authentication), data integrity (accuracy), and system-level defenses (architecture) - each of which surfaced independently in DIC stakeholder accounts.

3.1.4 Situating DIC against established governance frameworks

Internationally recognized frameworks such as the NIST Cybersecurity Framework and ISO/IEC 27001 organize data-security practice around functions that map closely onto what DIC stakeholders described in informal terms: identifying assets and risks, protecting them through access control and encryption, detecting anomalies, and maintaining an ongoing response and recovery capability. Similarly, principles embedded in data-protection regimes such as the EU's General Data Protection Regulation - purpose limitation, data minimization, and accountability for demonstrating compliance - echo the specific requests DIC respondents made spontaneously (non-disclosure to third parties, encryption, restricted access). None of these frameworks were named by respondents, and this paper does not claim DIC was consciously implementing them; rather, the alignment suggests that the underlying principles of sound data governance are intuitively recognized by non-specialist stakeholders even where the formal frameworks and legal instruments that codify those principles are absent from day-to-day institutional practice.

3.1.5 Cloud security posture and workload protection as governance analogues

A recent and directly relevant strand of work addresses how organizations formalize the kind of ad hoc access and privacy vigilance observed at DIC once systems scale into cloud infrastructure. Boamah's analysis of cloud security posture management (CSPM) shows that automated resource discovery, continuous compliance monitoring, and automated remediation can reduce misconfiguration incidents substantially and cut threat-detection time from weeks to hours relative to manual oversight [10]. In a healthcare-sector setting, Asante's framework for Cloud Workload Protection Platforms (CWPP) demonstrates how continuous visibility, runtime threat detection, and identity-centric controls can be organized into a coherent, workload-level governance architecture for protecting sensitive records processed in dynamic, containerized environments [11] - an architecture directly analogous to the ad hoc, individually-issued authentication and non-disclosure practices DIC respondents described. Separately, Asante's evaluation of confidential-computing runtimes shows that hardware-based Trusted Execution Environments can provide verifiable privacy guarantees and automated governance, risk, and compliance (GRC) enforcement even in multi-tenant, resource-shared environments, addressing exactly the blind spot DIC stakeholders flagged when they asked, without prompting, that their data be encrypted and not shared with third parties [12]. None of these frameworks were deployed at DIC; they are introduced here as the technical vocabulary that DIC's intuitive, informally-expressed governance instincts are already reaching toward.

3.1.6 The emerging role of artificial intelligence in compliance and access governance

A more recent strand of work, largely absent from the original thesis, concerns the use of AI in operationalizing data governance at scale. AI-assisted anomaly detection can flag unusual access patterns (e.g., a login from an unrecognized device or an atypical time of day) faster than manual audit processes allow, supporting the "regularly reviewing user access" practice that DIC respondents described as a manual, ad hoc task. Similarly, natural-language-processing tools are increasingly used to help small institutions map their internal data-handling practices against the requirements of privacy statutes such as Ghana's Act 843 or the EU's GDPR - a task that under-resourced institutions like DIC, which had no dedicated compliance officer among the stakeholder roles interviewed, would otherwise struggle to perform manually. Related work applying AI to cyber incident response demonstrates that automated, analytics-driven monitoring can meaningfully shorten the time between an anomalous event and its detection in operational settings [13], while separate work on usability-centred design for privacy-preserving configuration shows that non-expert users can be guided toward secure default behaviour when systems are designed with their cognitive and technical constraints in mind [14]. Neither line of work was developed for an institution with DIC's specific resource profile, but both suggest the underlying techniques are mature enough to be adapted downward. These applications remain nascent, and their reliability, especially for lower-resourced institutions with limited IT capacity, is not yet well established; this paper treats them as a research direction rather than a demonstrated solution (see Section 3.4).

3.1.7 Comparative reading against recent cloud-security case studies

It is instructive to place DIC's findings alongside the recent cloud-security literature reviewed above, since the underlying governance problem - how to give the right people access to sensitive data while keeping everyone else out, and how to prove that this is happening - is the same problem at very different scales. Boamah's CSPM analysis, conducted across enterprise multi-cloud deployments spanning AWS, Azure, and GCP, found that automated posture management reduced misconfiguration incidents by 60–80% and cut threat-detection time from weeks to hours relative to manual review [10]. DIC's "robust system" dimension - data accuracy, data security, and data privacy, all managed manually by a single IT associate - is, in effect, an unautomated analogue of exactly the posture-management problem CSPM tooling addresses at enterprise scale. The difference is not the nature of the problem but the resources available to solve it: DIC has no automated resource discovery, no continuous compliance monitoring, and no automated remediation, relying instead on a management member's one-time judgment that a vendor system had been "registered legally" before adoption. Similarly, Asante's CWPP framework for healthcare data warehouses [11] was designed explicitly because "traditional perimeter-centric security models are insufficient for modern healthcare cloud ecosystems" that continuously ingest and process protected information - a description that applies with only minor modification to a school system continuously ingesting attendance records, health information such as allergies, and academic histories for minors. The scale differs by orders of magnitude, but the underlying workload-protection logic - continuous visibility, identity-centric controls, and runtime monitoring rather than a one-time procurement check - is the same.

3.1.8 A threat-model reading of DIC's informal safeguards

Reframing the interview data as an informal threat model clarifies where DIC's practice is strong and where it is exposed. Identity spoofing risk is partially mitigated by unique logins and QR-code verification, but no respondent described a process for revoking access when a parent's custody status changes or a staff member leaves - an omission that, in more formal access-governance language, represents a missing de-provisioning control. Data-in-transit and data-at-rest exposure is addressed only at the level of stated expectation ("it should be encrypted"), not verified implementation; this is the single largest gap identified in this study, because an assumption of encryption that has never been audited provides no actual protection against interception or unauthorized database access. Third-party disclosure risk is addressed through a verbal norm communicated informally to at least one stakeholder, rather than a documented data-sharing agreement with the vendor operating the PROCARE platform - meaning the institution has limited ability to verify what the vendor itself does with the data once it leaves DIC's own systems. Insider-risk controls are the least developed dimension in the data: while "regularly reviewing user access" was named as a desirable practice by one facilitator, no respondent described an actual schedule, owner, or audit trail for such reviews.

3.1.9 Implications for policy and regulatory design

For Ghana's Data Protection Commission and comparable regulators elsewhere, the DIC case illustrates a specific and actionable enforcement gap: institutions are not resisting privacy obligations, they are simply operating without visibility into them. This suggests that awareness campaigns and low-cost compliance tooling aimed specifically at small private educational and community institutions - which typically lack a legal or compliance function entirely - may yield more practical improvement than additional statutory refinement of an already-adequate legal framework. A templated, sector-specific data-protection checklist for schools, distributed through existing regulatory or Ministry of Education channels, would map naturally onto the specific gaps this study identified: an access-provisioning and revocation procedure, a documented encryption standard, and a written third-party data-sharing policy.

DIC's Informal Practice Mapped to Formal Governance Concepts

Formal concept	DIC's informal analogue
Identity and access management	Unique login credentials; QR-code pick-up verification
Data-in-transit / at-rest encryption	Verbal expectation, unverified implementation
Third-party data-sharing agreement	Informal verbal norm relayed by an administrator
Access review / de-provisioning	Named as desirable, no schedule or owner identified
Regulatory compliance mapping (Act 843)	Not referenced by any respondent, including IT and admin roles
Continuous compliance monitoring (CSPM)	One-time procurement-stage legitimacy check
Workload-level protection (CWPP)	No workload-level distinction; system treated as a single unit

3.2 Findings

Stakeholders consistently described authentication as the mechanism through which privacy and security expectations were operationalized in practice, even without formal reference to Ghanaian data-protection law. A parent explained receiving "a unique code to my email to ensure security" from the system administrator, functioning as an ad hoc authorization step. Another parent described the PROCARE app's QR-code check-in as a way the "system recognizes me as an authorized pick-up person" - an identity-verification control adopted for child-safety reasons that doubles as an access-control mechanism. A procurement officer reported being advised that "log in details and unique number of users is not shared to any third party," an informal non-disclosure norm functioning in place of a documented data-sharing policy.

Security and privacy language recurred independently across roles, as summarized in Table 2. An IT associate referenced "privacy and security features of the system" as a distinct consideration from system functionality. A facilitator described security "in terms of data not attaining by third parties." A teacher raised "data privacy" and "privacy of data" without prompting about specific legislation, and another facilitator noted that "students' details are kept confidential." A parent tied these controls directly to safety: "I do not want my child's information and my information out there…every user has a unique login detail…my digital signature is even required." Administrators and management, for their part, connected system trustworthiness to legal legitimacy, with one management member noting that before adoption, "we examined how safe it was for the school and if the system has been registered legally…to show patent right," and a procurement officer stating that "the system should be rigorously validated to ensure its accuracy and reliability." Encryption was explicitly requested by at least one respondent, who stated that "the data stored in PROCARE should be encrypted to protect it from unauthorized access" - a specific technical control named by a non-technical stakeholder, indicating that privacy expectations at DIC exceeded what the institution's actual data-handling practice appeared to guarantee.

Illustrative Stakeholder Statements on Access, Authentication, and Privacy

Role	Illustrative statement
Parent	"Administrator sent me a unique code to my email to ensure security."
Parent	"The system recognizes me as an authorized pick-up person."
Procurement officer	"Log in details…not shared to any third party."
IT associate	"Privacy and security features of the system."
Facilitator	"Security in terms of data not attaining by third parties."
Teacher	"Data privacy" / "Privacy of data."
Facilitator	"Students' details are kept confidential."
Parent	"The data stored in PROCARE should be encrypted."
Management member	"We examined how safe it was…if the system has been registered legally."
Procurement officer	"The system should be rigorously validated to ensure its accuracy and reliability."

System design and validation surfaced as a related but analytically distinct concern from authentication. An administrator stated plainly that "the system should be designed to ensure the confidentiality and security" of the

records it holds, treating design-stage decisions as a security matter separate from operational access controls. A teacher raised the practical corollary of design-stage assurance - ongoing maintenance - noting the importance of "regularly backing up data," which, while framed by the respondent as a continuity concern, is equally a resilience control against data loss from compromise or hardware failure. A parent connected system legitimacy to intellectual-property and registration status, asking "how the system was built and if the patent right" had been secured, indicating that at least some stakeholders treated legal registration as a proxy for trustworthiness in the absence of any technical means of verifying the vendor's security practices directly. A procurement officer, finally, insisted that "PROCARE should be used in a non-discriminatory manner," extending the governance conversation beyond confidentiality and integrity to encompass equitable access as a further, if less technical, governance concern.

Taken together, these findings show privacy and security operating as felt obligations enforced through informal institutional norms and individual vigilance, rather than through reference to Ghana's Data Protection Act or Public Records and Archival Administration Act. No respondent - including administrative and IT roles - referenced either statute by name.

3.2.1 Ghana's data-protection framework in comparative regional context

Ghana's Data Protection Act, 2012 (Act 843) was among the earlier comprehensive data-protection statutes on the African continent, predating Nigeria's Data Protection Act, 2023 and following a similar principles-based structure to South Africa's Protection of Personal Information Act (POPIA), 2013 - both of which draw, in turn, on the EU's data-protection tradition that culminated in the GDPR. In principle, Act 843 requires data controllers to register with the Data Protection Commission, obtain consent for processing, and implement appropriate security safeguards; a school information system processing minors' academic, attendance, and health-adjacent data would, on a plain reading, fall squarely within its scope. What this case study adds to that regional legislative picture is not a claim about the adequacy of the statute itself, but an empirical observation about reach: a comprehensive, GDPR-adjacent legal framework existing on the books does not, by itself, guarantee that the institutions it is meant to govern are aware of it, let alone actively compliant with it. This distinction - between legislative adequacy and operational reach - is easy to overlook in comparative data-protection scholarship that focuses primarily on the content of statutes rather than on ground-level institutional awareness, and this case study suggests the latter deserves at least as much research attention as the former, particularly for small, non-public institutions that fall outside the large-enterprise compliance apparatus that most data-protection enforcement activity in practice tends to target.

4 DISCUSSION

These findings support the authentication-as-privacy-protection literature [5,6] at the level of lived institutional experience: DIC stakeholders treated identity verification (login credentials, QR codes, digital signatures) as functionally inseparable from privacy protection, describing both in the same breath. But the findings also expose a governance gap consistent with Tsvuura and Ngulube [2] and Asogwa [3]: none of the safeguards described by respondents were traceable to a specific legal requirement. They were institutional improvisations — sensible, but not standardized, audited, or legally anchored.

This has two practical implications. First, encryption, access logging, and third-party data-sharing restrictions were requested or assumed by stakeholders but not confirmed as implemented; the gap between stated expectation and verified practice is itself a security risk, since unverified assumptions about protection ("it should be encrypted") are not the same as an audited encryption policy. Second, the absence of reference to Act 843 or Act 535 in interviews conducted with administrative and IT-facing roles suggests that Ghana's existing data-protection framework, however sound on paper, is not reaching the operational level of small private institutions - consistent with Asogwa's [3] and Mukred et al.'s [4] broader observation about enforcement gaps in Sub-Saharan Africa. Read against the cloud-security literature discussed in Section 3.1, DIC's informally expressed governance instincts (identity-centric access, non-disclosure, encryption) are precisely the properties that CSPM and CWPP architectures are designed to enforce systematically [10,11] - suggesting that the gap at DIC is one of tooling and enforcement capacity rather than conceptual misunderstanding of what good governance requires.

4.1 Cross-case comparison with other Sub-Saharan African settings

The pattern observed at DIC is consistent with, and adds granularity to, prior cross-national findings. Tsvuura and Ngulube's [2] study of two Zimbabwean universities found record-keeping legislation lagging behind technological change at the level of national policy; DIC's data show what that lag looks like from the vantage point of a single institution trying to operate a modern information system underneath an unenforced legal umbrella. Where the Zimbabwean study observed the legislative gap directly, DIC's interview data reveal the downstream behavioral adaptation: institutional actors do not wait for enforcement, they independently reconstruct plausible-sounding governance practices (unique logins, verbal non-disclosure norms, a one-time legitimacy check at procurement) that

approximate, without ever citing, what formal legislation would otherwise require. Asogwa's [3] and Mukred et al.'s [4] broader observation that most Sub-Saharan African countries lack dedicated electronic-records legislation is, in Ghana's case, not quite accurate at the level of the statute book - Act 843 exists and is reasonably comprehensive - but is empirically accurate at the level of institutional awareness and operational practice, which is arguably the more consequential gap for practitioners and regulators alike.

4.2 Why informal governance is not sufficient, even when it happens to work

It is tempting to read DIC's findings optimistically: stakeholders behaved responsibly, requested reasonable safeguards, and the institution suffered no reported breach during the study period. This reading should be resisted for two reasons specific to the evidence. First, absence of a reported breach is not evidence of absence of vulnerability; none of the technical controls stakeholders assumed were in place (encryption, access logging, revocation procedures) were independently verified as part of this study, so the finding is about stated expectation, not confirmed protection. Second, informal governance is inherently non-transferable: it depends on the specific individuals currently holding institutional knowledge (the single IT associate, the management member who vetted the vendor at procurement) rather than on a documented, auditable process that would survive staff turnover. A governance arrangement that cannot survive the departure of one employee is not a resilient arrangement, regardless of how sensible it looks at a single point in time.

4.3 Future Research Directions

Given the resource constraints evident in this case - a single IT associate managing a multi-stakeholder system, no named compliance role, and reliance on informal norms — future research should examine whether AI-assisted tools can close the gap between stated privacy expectations and verifiable practice without requiring institutions to hire dedicated compliance staff. Four directions are proposed. First, AI-based anomaly detection for access logs, to operationalize the "regularly reviewing user access" practice that respondents described as manual and inconsistent, in the spirit of the continuous compliance monitoring shown to reduce misconfiguration incidents in cloud environments [10]. Second, automated, AI-assisted mapping of an institution's actual data-handling practices against Act 843 and comparable frameworks, producing a gap report a non-specialist administrator could act on. Third, natural-language consent and privacy-notice generation tools, so that small institutions can produce legally coherent privacy policies without legal counsel. Fourth, confidential-computing-based access enforcement, adapted from multi-tenant cloud settings [12], could in principle provide verifiable technical guarantees - rather than informal assurances - that data is not exposed to unauthorized parties, directly answering the encryption and non-disclosure expectations DIC stakeholders raised unprompted. Each direction requires empirical validation - particularly around the accuracy and false-positive rates of AI anomaly detection in low-connectivity environments - before it can be recommended as institutional practice, and that validation is left to future work.

4.3.1 Sequencing the research agenda

The four directions proposed above are not equally tractable in the near term, and a sensible research agenda would sequence them by feasibility rather than pursue all four simultaneously. Automated statutory gap-mapping (the second direction) is arguably the most immediately tractable, since it requires only natural-language processing over a relatively small, stable body of text (Act 843 and comparable regulations) matched against an institution's self-reported practices gathered through a structured questionnaire similar to the one used in the parent study; no novel sensor data or system integration is required. AI-based anomaly detection for access logs (the first direction) is more technically demanding, since it requires the target institution to already be generating structured, timestamped access logs - a precondition DIC does not currently meet, since access appears to be tracked informally rather than logged systematically. Confidential-computing-based enforcement (the fourth direction) is the least tractable in the near term for an institution like DIC, since it presumes a level of infrastructure investment and technical sophistication - specialized hardware, integration with a Trusted Execution Environment - that is arguably a poor fit for the same resource-constrained institutions this paper is centrally concerned with; it may be more relevant to the vendor operating PROCARE than to DIC itself, since the vendor operates at the multi-tenant scale where confidential computing's guarantees become most valuable. Natural-language consent and privacy-notice generation (the third direction) sits between these poles in difficulty and could plausibly be piloted at DIC within a single academic year, making it a reasonable starting point for a follow-on empirical study.

4.3.2 A note on vendor responsibility

A theme that emerges from reading the four future-research directions together is that several of the more resource-intensive controls - workload-level protection, confidential computing, and continuous compliance monitoring in particular - are more naturally the responsibility of the vendor operating PROCARE (or a comparable school information system) than of any individual client institution. DIC, as a single school with one IT associate, is poorly positioned to independently implement CWPP-style workload protection [11] or CSPM-style posture management [10] across its use of a third-party platform. This suggests that future research and, more directly, future procurement practice, should

treat vendor-side security posture as a first-order criterion in system selection, rather than an afterthought addressed only if a stakeholder happens to ask about it - as appears to have been the case at DIC, where system legitimacy was assessed once, informally, at the point of adoption.

4.4 Practical Implications for Practitioners

For school administrators and similarly resourced institutional leaders, three implications follow directly from these findings. First, the gap between what users assume is happening to their data ("it should be encrypted") and what is actually documented as institutional policy should be treated as an audit priority, not a reassurance to be accepted at face value; an institution that has never formally verified its encryption-at-rest or in-transit posture cannot know whether it matches user expectation. Second, the ad hoc, individually-issued authentication practices described by respondents (a unique code emailed by an administrator, a QR code tied to a single authorized pick-up person) work at small scale but do not obviously scale to institutions with larger enrollment or higher staff turnover; a written access-provisioning and de-provisioning procedure - who is granted access, when, and how it is revoked when a role changes - would formalize what DIC currently manages informally. Third, given that no respondent referenced Ghana's Data Protection Act or the Public Records and Archival Administration Act by name, institutions should not assume staff or management are aware of their statutory obligations simply because they are behaving in privacy-conscious ways; awareness of the law and intuitive privacy-protective behavior are evidently two separate things at DIC, and only the latter was observed.

4.5 Toward a Minimum Viable Governance Checklist

Drawing directly on the gaps identified in Sections IV through VII, Table 3 proposes a minimum viable governance checklist intended for institutions with DIC's resource profile: a single IT staff member or equivalent, no dedicated compliance function, and a vendor-supplied information system rather than an in-house build. The checklist is deliberately narrow - it does not attempt to replicate enterprise-grade governance frameworks such as ISO/IEC 27001 in full, since the evidence from this case suggests that a small institution asked to adopt a comprehensive enterprise framework wholesale is unlikely to sustain it. Instead, each item corresponds to a specific gap directly observed in the interview data, on the theory that closing the exact gaps found is more achievable, and more likely to be sustained, than pursuing general best practice in the abstract.

Beyond the six items listed in Table 3, institutions adopting this checklist should treat it as a starting baseline rather than a ceiling. A natural extension, once the six baseline items are in place and sustained for a full academic year, would be to introduce lightweight logging of access events - even a simple spreadsheet recording who accessed which category of record and when - as a precondition for the AI-assisted anomaly detection direction proposed in Section 3.4. Without such logging, there is no data for an anomaly-detection system to analyze, however sophisticated the underlying model; this is a common and underappreciated failure mode in discussions of AI-assisted security tooling, where the promise of the analytic method obscures the more basic prerequisite of having structured data for it to operate on in the first place.

It is also worth noting what the checklist deliberately omits. It does not include a formal Data Protection Officer role, since institutions of DIC's size are unlikely to be able to justify a dedicated position; instead, the checklist assigns each responsibility to an existing role (the IT associate, a management member, an administrator) rather than creating a new one. It does not include a formal incident-response plan, on the grounds that an institution that cannot yet reliably answer "who has access to what" is not yet ready to benefit from a formal plan for responding to unauthorized access - the foundational visibility work should come first. And it does not include specific technology recommendations (which encryption standard, which vendor), since the evidence from this study concerns institutional process and awareness rather than technology selection, and specific technology recommendations would in any case date quickly relative to the underlying governance principles the checklist is meant to encode.

Each of these six actions is deliberately achievable without new hires or significant budget, addressing the criticism, raised in Section 3.5, that institutions in DIC's position cannot simply be told to adopt enterprise-grade governance. Whether such a minimal checklist meaningfully reduces risk relative to the status quo - and whether institutions that adopt it sustain the practice over time - is itself an empirical question this paper does not resolve, and is proposed as a natural extension of the AI-assisted tooling research agenda outlined in Section 3.4: an AI-assisted assistant could, for instance, walk a non-specialist administrator through exactly these six items on a recurring basis, lowering the effort required to sustain them.

4.5.1 Minimum Viable Governance Checklist Derived from DIC Findings

Gap observed	Minimum action
No de-provisioning process	Written procedure: who removes access, and within what time frame, when a role or custody status changes
Unverified encryption	One-time technical audit confirming encryption at rest and in transit; document the result
Verbal-only non-disclosure norm	One-page data-sharing policy signed by vendor and institution
No access-review schedule	Quarterly access-list review assigned to a named owner
No statutory awareness	One-hour briefing for administrative staff on Act 843 obligations
No vendor security visibility	Annual request for vendor's own security and compliance attestation

Each of these six actions is deliberately achievable without new hires or significant budget, addressing the criticism, raised in Section 3.5, that institutions in DIC's position cannot simply be told to adopt enterprise-grade governance. Whether such a minimal checklist meaningfully reduces risk relative to the status quo - and whether institutions that adopt it sustain the practice over time - is itself an empirical question this paper does not resolve, and is proposed as a natural extension of the AI-assisted tooling research agenda outlined in Section 3.4: an AI-assisted assistant could, for instance, walk a non-specialist administrator through exactly these six items on a recurring basis, lowering the effort required to sustain them.

Fourth, and more broadly, institutional leaders should recognize that the burden of privacy protection at DIC currently falls disproportionately on informal vigilance by individual staff and parents rather than on documented institutional process - the procurement officer who was told logins are not shared, the parent who requested encryption, the management member who personally vetted the vendor's legitimacy. This distributed, person-dependent model of governance is fragile precisely because it is not written down anywhere; if any of these specific individuals left the institution, the knowledge and the practice attached to it would likely leave with them. Converting person-dependent norms into institution-dependent documentation - however minimal - is therefore not merely a compliance nicety but a basic continuity safeguard.

limitations

This paper reframes findings from a single-institution, qualitative case study; the authentication and privacy practices described at DIC - a private school with a specific stakeholder mix of parents, teachers, and a small administrative and IT staff - may not generalize to public institutions, larger enrollment settings, or non-educational record-holding organizations. The interviews probed user experience of the PROCARE system generally, not privacy law compliance specifically, so the finding that no respondent referenced Act 843 or Act 535 should be read as suggestive rather than a rigorously measured awareness gap; a study designed explicitly to test statutory awareness among administrative staff would be needed to confirm this pattern. As is common with cross-disciplinary theses, a small number of sources cited in the literature review of the parent study (Ahmad et al., 2023; Cheng et al., 2017; Telo, 2021) are reproduced here with the year and general finding as reported in that source.

A further limitation concerns the comparative material introduced in Sections II-E and V-A: the cloud security posture management, workload protection, and confidential computing literature cited there was not part of the original case-study design and was not tested against DIC's actual infrastructure. These comparisons are offered as a conceptual bridge between an informally-governed small institution and the formal governance architectures developed for much larger, cloud-native environments, not as evidence that any specific tool or framework from that literature would perform as described if deployed at an institution with DIC's connectivity, budget, and staffing profile - a profile substantially different from the enterprise and healthcare settings in which that literature's findings were originally generated. Finally, this paper's re-analysis was conducted on already-collected, already-coded data rather than through new fieldwork; a study designed from the outset around the governance and security questions this paper raises might elicit richer or different findings than a re-reading of data collected for a broader digital-preservation study could surface.

4.5.2 Toward a research agenda for institutional data governance in low-resource settings

Beyond the specific directions already proposed, this case study points to a broader research agenda that the digital preservation and information systems literatures have only begun to address: how should governance frameworks designed for enterprise, cloud-native, or healthcare-scale environments be adapted - rather than simply scaled down -

for institutions that will never have a dedicated compliance function, a security operations center, or a budget for commercial CSPM tooling? Simply scaling down an enterprise framework by removing steps until it fits a small institution's resources risks removing exactly the steps that made the framework effective in the first place. A more productive research direction may be to start from the resource constraint itself - one IT staff member, intermittent connectivity, no legal counsel - and ask what the minimum sufficient governance architecture looks like when designed for that constraint from first principles, rather than as a subtraction from an enterprise ideal. The minimum viable checklist proposed in Section 3.6 is offered as a first, modest attempt at this kind of resource-first design, not as a finished answer.

5 CONCLUSION

Divine International College's experience shows that user authentication is functioning as an informal privacy safeguard in a Ghanaian school information system, ahead of, rather than because of, formal legal enforcement. Institutions in similar resource-constrained, digitizing-but-under-regulated environments should not wait for stronger legal enforcement to formalize what their users already expect: documented encryption standards, an explicit data-sharing policy, and periodic access-authorization reviews. Regulators, in turn, should treat the gap between Act 843's requirements and its operational visibility at the institutional level as an enforcement and awareness priority, not merely a legislative one. Future AI-assisted compliance tooling, cloud security posture management, and confidential-computing-based access enforcement may help close this gap for institutions that cannot otherwise afford dedicated data-governance staff, but such tools should be evaluated rather than assumed to work in the resource-constrained settings where they are most needed.

More broadly, this case study suggests that the relationship between formal data-protection legislation and actual institutional practice, at least in resource-constrained settings, is far weaker than either policymakers or technology vendors typically assume. Ghana's Data Protection Act, 2012 is a reasonably comprehensive piece of legislation by regional standards, yet it exerted no detectable influence on how DIC's stakeholders described, understood, or implemented the safeguards they nonetheless valued. This is not evidence that the law has failed; it is evidence that the law's existence and its operational reach are two separate achievements, and that closing the distance between them may depend less on further legislative refinement than on low-cost, practically deployable tooling - whether a simple paper checklist of the kind proposed in Section 3.6, or the AI-assisted compliance and posture-management approaches discussed in Sections II and VI - that meets institutions where they actually are, rather than where policy assumes them to be.

Finally, it is worth stating plainly what this paper does and does not claim. It does not claim that DIC is unusually poorly governed; the evidence, if anything, suggests the opposite - stakeholders across every role interviewed showed a consistent, unprompted instinct toward privacy-protective behavior. What the paper claims is narrower and, in some ways, more concerning: that even reasonably conscientious institutional behavior, absent documented process and independent verification, leaves privacy and security resting on the memory and goodwill of specific individuals rather than on an institutional structure that would survive their departure. That distinction - between good instincts and durable governance - is the central finding this paper contributes to the literature on digital preservation, data privacy, and institutional security in resource-constrained settings, and it is the distinction future AI-assisted and vendor-side tooling research, outlined in Section 3.4, should be designed to close.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Acknowledgments

The author thanks the stakeholders of the case-study institution for their participation in the interviews on which this study draws, and the administrative staff who facilitated access for the original data collection.

Disclosure of conflict of interest

The author declares no financial or non-financial conflict of interest related to this work.

Statement of ethical approval

The interview data underlying this study were collected as part of a broader academic research project examining digital preservation and data management at the case institution. Ethical clearance and institutional permission for the original data collection were obtained through the supervising academic institution prior to fieldwork, and all participants were briefed on the purpose of the research before interviews were conducted.

Statement of informed consent

Informed consent was obtained from all individual participants included in the study. Participants were informed of the voluntary nature of their participation and their right to withdraw at any time.

REFERENCES

[1] Marciano R, Lemieux V, Hedges M, Esteva M, Underwood W, Kurtz M, Conrad M. Archival records and training in the age of big data. Adv Librariansh. 2018;44B:179-199.

[2] Tsvuura G, Ngulube P. Digitisation of records and archives at two selected state universities in Zimbabwe. J South Afr Soc Arch. 2020;53:20-34.

[3] Asogwa BE. The challenge of managing electronic records in developing countries: implications for records managers in sub Saharan Africa. Rec Manag J. 2012;22(3):198-211.

[4] Mukred M, Yusof ZM, Mokhtar UA, Manap N. Electronic records management system adoption readiness framework for higher professional education institutions in Yemen. Int J Adv Sci Eng Inf Technol. 2019;9(2):464-473.

[5] Hasal M, Nowaková J, Ahmed Saghair K, Abdulla H, Snášel V, Ogiela L. Chatbots: security, privacy, data protection, and social aspects. Concurr Comput Pract Exp. 2021;33(19):e6426.

[6] Ahmad S, et al. Multi-factor authentication and biometric access control for sensitive data protection [full venue details not independently verified; author to confirm before submission]. 2023.

[7] Dika A, Nowostawski M. Security vulnerabilities in Ethereum smart contracts. In: Proceedings of the IEEE International Conference on Internet of Things (iThings), IEEE GreenCom, IEEE CPSCom, and IEEE SmartData; 2018. p. 955-962.

[8] Telo J. System-level security architecture: firewalls, intrusion detection, and audit practice [full venue details not independently verified; author to confirm before submission]. 2021.

[9] Cheng L, et al. Security auditing practices in enterprise information systems [full venue details not independently verified; author to confirm before submission]. 2017.

[10] Boamah KG. Cloud security posture management: a comprehensive analysis of automated risk identification and mitigation in multi-cloud environments. Int J Res Innov Soc Sci. 2025;9(11):4458-4471.

[11] Asante A. Strengthening healthcare cloud security using Cloud Workload Protection Platforms (CWPP): a framework for protecting patient-critical workloads in health data warehouses. Int J Res Innov Soc Sci. 2025;9(11):7863-7889.

[12] Asante A. Evaluating confidential computing runtimes to enforce verifiable privacy guarantees and automated GRC controls in multi-tenant cloud data processing workflows. Int J Comput Appl Technol Res. 2025;14(11):40-48.

[13] Boamah KG, Asante A, Timean A, Okai KF. Artificial intelligence integration in cyber incident response teams to enable faster containment, forensic accuracy, and resilient business continuity. Int J Sci Res Arch. 2025;17(1):1263-1280.

[14] Boamah K. Usability standards for privacy-preserving security configuration in IoT devices. Int J Res Publ Rev. 2025;6(11).

[15] Republic of Ghana. Public Records and Archival Administration Act, 1997 (Act 535). Accra: Government of Ghana; 1997.

[16] Republic of Ghana. Data Protection Act, 2012 (Act 843). Accra: Government of Ghana; 2012.