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

INFORMATION LITERACY AS A HUMAN SECURITY LAYER: BRIDGING DIGITAL COMPETENCY GAPS TO REDUCE DATA-HANDLING RISK IN SCHOOL INFORMATION SYSTEMS

Kwaku Gyamfi Boamah and Afua Asantewaa Asante *

College of Computing – Cybersecurity Grand Valley State University, USA.
* Corresponding Author

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

International Journal of Science and Research Archive, 2026, 20(02), 323–334

Publication history: Received on 21 June 2026; revised on 12 August 2026; accepted on 14 August 2026

Article DOI: https://doi.org/10.30574/ijsra.2026.20.2.1559

ABSTRACT

Data breaches and privacy failures are frequently attributed to weak technical controls, yet a large share of institutional risk originates upstream, in how confidently and correctly staff and users operate the systems that hold sensitive data. This study draws on qualitative case-study evidence from a Ghanaian school using a school management information system to examine information literacy as a human security layer rather than a purely operational competency. Interviews with 25 stakeholders reveal that user unfamiliarity with the system, including difficulty with navigation, resistance to change, and reliance on legacy paper-based habits, was widespread during and after adoption, even among staff directly responsible for records and data entry. The findings are discussed in relation to literature on information literacy, human-factors security, and technology-adoption curves, and situated within a security lens: low system literacy does not only slow adoption, it creates the conditions for data-handling errors, workaround behaviours, and inconsistent application of privacy and security practices that a system's technical controls cannot compensate for on their own. The study closes by outlining directions for artificial-intelligence-assisted, adaptive literacy training, drawing on recent work in usability-centred design for privacy-preserving configuration, as an underexplored way to close competency gaps in resource-constrained institutions without dedicated training budgets.

Keywords: Information Literacy; Human Factors; Data Security; Digital Preservation; Artificial Intelligence

1 INTRODUCTION

Much of the discourse on data security in digitizing institutions focuses on firewalls, encryption, and access control - technical layers that assume a competent, confident user on

the other end of the login screen. That assumption often does not hold. In institutions adopting new information systems for the first time, the people operating the system are frequently learning it under pressure, alongside their existing responsibilities, with limited formal training. This paper argues that information literacy - the ability to find, evaluate,

and correctly use information and information systems - functions as a human security layer: a competency gap here does not simply slow down operations, it creates a class of risk that technical controls cannot fully offset.

This paper makes four contributions. First, it provides empirical evidence, drawn from a live institutional rollout, that information-literacy gaps manifest as concrete, security-relevant workaround behaviors (specifically, parallel paper record-keeping) rather than merely as slower task completion. Second, it shows that these gaps are not confined to any single stakeholder category - they appear among parents, teachers, and IT-facing staff alike - challenging any assumption that literacy risk can be addressed by training only the least technical users. Third, it connects this small-institution finding to recent work on usability-driven security design, arguing that the same usability principles developed for IoT device configuration apply directly to institutional information-system onboarding. Fourth, it proposes a sequenced, resource-appropriate research agenda for AI-assisted adaptive literacy support, explicitly designed for institutions without dedicated training budgets.

This paper revisits data originally collected as part of a broader study of digital preservation and data management at Divine International College (DIC), Accra, isolating the information-literacy dimension of that study and reframing it through a security and data-handling lens, in dialogue with human-factors and information-literacy literature. It is worth being explicit about scope at the outset: this paper is not a general argument that more training is always beneficial, nor a claim that DIC's staff or parents were poorly prepared in some deficient sense. Rather, it is a narrower, more specific argument - that a predictable, temporary competence gap during any new system's adoption has an identifiable and measurable security signature, and that this signature deserves dedicated research and institutional attention distinct from generic discussions of "user training" or "change management," both of which tend to treat the transition period as a productivity problem rather than a security one. The remainder of the paper proceeds as follows. Section 3.1 reviews literature on information literacy, human-factors security, technology adoption, and usability-driven security design. Section 2 describes the qualitative methodology. Section 3.2 presents findings. Section 3.3 discusses their implications. Section 3.4 proposes an AI-assisted research agenda. Section 3.5 offers practical recommendations, Section 3.6 proposes a minimum viable onboarding checklist, Section 3.7 discusses limitations, and Section 3.8 concludes.

2 MATERIAL AND METHODS

2.1 Research paradigm

This paper draws on the same interpretivist and constructivist qualitative paradigm described in the companion governance analysis [4]. Interpretivism holds that human behavior is too complex to study through probabilistic models and that knowledge is created by interpreting the meanings people attach to events; constructivism holds that individuals construct their own understanding of experience through reflection. Both were judged appropriate here because the phenomenon under study - how stakeholders experienced learning and adapting to a new information system - is a matter of subjective interpretation rather than an objectively measurable quantity.

2.2 Case, sampling, and data collection

This paper draws on the same qualitative case-study dataset described in the companion analysis of DIC's data-governance practices: 25 of 27 targeted stakeholders interviewed (83.3%), across parent, facilitator, teacher, administrator, IT, procurement, and management roles, selected via a maximal-variation, purposive and convenient sampling strategy and coded using a first-order/second-order/aggregate-dimension approach. This paper isolates two aggregate dimensions produced by that coding - "system usage" (verification, communication, data management, assessment, records keeping) and "conservative users" (system navigation, mindset, adaptation, incompatible devices, manual record keeping) - for focused discussion.

2.3 Reliability, validity, and reflexivity

The same reliability and validity procedures described in the companion paper apply to this analysis: a supervisor-reviewed interview guide, advance notice to respondents, systematic coding review, a pilot round to check construct validity, and interview transcripts returned to respondents for member-checking to support internal validity. As with the companion analysis, the researchers had an existing relationship with the case institution, which likely eased candor about difficulties and failures but means the analytic categories applied to the data reflect the research team's own coding scheme rather than categories stakeholders would necessarily have generated unprompted.

2.4 Reflexivity and interview protocol

As in the companion governance analysis [4], the interview protocol underlying this dataset used open-ended rather than checklist-style questions, asking respondents to describe their experience learning and using the system in their own words rather than rating their proficiency against a predetermined scale. This design choice is likely responsible

for the richness of the workaround-behavior data reported in Section 3.2: a checklist question ("Did you experience difficulty? Yes/No") would likely have understated the specific, security-relevant detail - exactly which tasks prompted a reversion to paper, and for how long - that open-ended description surfaced. The researchers' existing relationship with the case institution, noted in the companion paper, applies equally here and should be read as a factor that likely increased candor about difficulty and failure, while also meaning the coding categories applied to the data ("system usage," "conservative users") reflect the research team's own analytic framing rather than categories respondents generated unprompted. Member-checking - returning transcripts to respondents for confirmation before analysis - was used specifically to guard against this risk.

3 RESULTS

3.1 Literature Context

3.1.1 Information literacy and data management

Effective data management depends on more than system design; it depends on the competence of the people who use the system daily. Information-literacy scholarship argues that efficient data management within an organization necessarily depends on the information literacy skills of both information professionals and end users [1]. Related work finds that individuals with strong information literacy skills are more likely to engage in continuous learning and adapt to evolving data management practices [2,3] - staying current with tools, methods, and standards rather than falling back on ad hoc or outdated habits when a new system is introduced.

3.1.2 Information literacy as a security variable

Human-factors research in information security has long recognized that misconfiguration, inconsistent practice, and workaround behavior by well-meaning users are a leading cause of data exposure - often exceeding the risk posed by deliberate external attack. A user who does not understand why a control exists (for example, why a unique login should not be shared, or why a record should not be exported informally) is more likely to bypass it under time pressure, not out of malice but out of unfamiliarity. This reframes information literacy from an operational-efficiency concern into a security variable: every stakeholder who struggles to navigate a system, or who reverts to a parallel paper process out of discomfort with the digital one, represents a point where the institution's actual data-handling practice diverges from its intended one.

3.1.3 The limits of purely technical security investment

This has a direct implication for how institutions allocate scarce security investment. A well-designed authentication system - of the kind examined in a companion analysis of DIC's authentication practices [4] - can be undermined if the people using it do not understand or trust it enough to use it consistently. Conversely, institutions with modest technical budgets may gain more security benefit, per dollar spent, from structured onboarding and literacy support than from additional technical controls layered onto a system that staff and parents do not yet use confidently.

3.1.4 Adoption curves and the security-relevant transition window

Technology-adoption literature has long recognized that new-system rollout follows a predictable curve, with early resistance and errors giving way to competence and, eventually, confident use. What the DIC findings add to this general pattern is a security-specific observation: the transition window is not merely a productivity dip, it is a period during which the institution's actual data-handling practice - as opposed to its intended, system-mediated practice - is least predictable and least auditable, because some information is moving through informal, undocumented paper channels precisely while formal digital channels are being learned. This suggests that the length of an institution's adoption window is itself a security-relevant variable worth measuring and, where possible, shortening.

3.1.5 Usability as a precondition for secure behavior: lessons from IoT

A directly relevant parallel comes from recent work on usability-driven privacy and security configuration in consumer IoT devices. Boamah's analysis of usability standards for privacy-preserving IoT configuration argues that complex or poorly designed configuration interfaces frequently lead to suboptimal security settings, user misconfiguration, or complete neglect of privacy controls, and proposes human–computer-interaction usability standards - incorporating privacy heuristics such as minimization, transparency, and informed consent - specifically designed to guide non-expert users toward secure behavior without imposing excessive cognitive burden [5]. The parallel to DIC is direct: PROCARE, like the IoT devices examined in that work, is operated by non-expert users (parents, teachers) rather than IT professionals, and the interview data show precisely the pattern that usability-centered design is meant to prevent - users falling back on manual workarounds not because they reject the system's security goals, but because the interface and onboarding process did not make secure, confident use achievable within the cognitive and time budget available

to them. This suggests that the same evaluation metrics proposed for IoT usability - configuration accuracy rates, error reduction, user decision confidence, adherence to recommended baselines [5] - could be adapted as a diagnostic instrument for institutional information-system onboarding, offering a measurable way to detect the literacy gap this paper documents qualitatively, before it manifests as a workaround behavior with security consequences.

3.1.6 Artificial intelligence and adaptive literacy training

A largely unexplored direction, absent from the original case study, is the use of AI-assisted, adaptive training tools to close information-literacy gaps at low marginal cost. Rather than a single generic onboarding session, adaptive systems can tailor guidance to a user's demonstrated proficiency - for instance, providing extra walkthroughs for the "system navigation" and "incompatible device" difficulties that recur in this dataset, while moving more confident users through faster. AI-based chat assistants embedded in school information systems could, in principle, answer a parent's or teacher's in-the-moment question about how to complete an action correctly, reducing the improvisation and workaround behavior that create data-handling risk. As with the compliance-tooling directions discussed in the companion governance paper [4], these applications are proposed here as a future-research agenda, not a validated intervention; their effectiveness for older or lower-digital-literacy users, and their cost relative to simpler in-person training, remains to be tested.

3.2 Information literacy in the Ghanaian educational technology context

The literacy gap documented in this case study should also be read against the broader context of educational technology adoption in Ghana and comparable Sub-Saharan African settings, where school-level digitization has often proceeded faster than the training infrastructure needed to support it. Government and donor-funded digitization initiatives in the region have historically prioritized hardware and connectivity provision - the infrastructure concerns examined in a companion analysis [6] - over sustained, role-specific training for the staff and parents who must actually operate the resulting systems day to day. DIC's experience, as a privately funded institution making its own procurement and training decisions rather than following a donor-prescribed rollout plan, suggests that this training gap is not solely a function of donor program design; it recurred even where the institution had full control over its own adoption process, suggesting a more general pattern in which literacy support is treated as a lower priority than the underlying technology procurement, regardless of who is making the purchasing decision.

Findings

Stakeholders across roles described genuine day-to-day reliance on the system for core tasks: a facilitator explained using PROCARE "to check attendance, who is dropping off and picking up students," while an administrator described using it "to send and receive information from parents." An IT associate described using the system "for preparing invoice of student and managing their portfolio," and a management member described a broader oversight role: "I oversee the general operation of the system and how it is used by users. I give access to users, input user details, and generate operational reports."

Alongside this competent, routine use, a distinct and equally consistent pattern of difficulty emerged, summarized in Table 1. A parent described early trouble: "I encountered challenge with the network and the QR scan code. Also, my phone at that time could not support the system." A procurement officer reported ongoing friction: "it requires me to be on the phone to check for notification most of the time." A homeroom facilitator admitted, "I encountered a little bit of difficulty as a result of change of system," while another facilitator was more direct: "I was still used to the traditional way of keeping record, which made it difficult for me to use PROCARE." An administrator acknowledged, "it took time to transition," and a teacher recalled, "Yes, we did [face difficulty], as it was a new environment." An IT administrator - a role one would expect to be most system-confident - nonetheless reported: "Yes, I did. It was all new to me so getting used to it in the early stages was extremely difficult." A management member summarized the institutional cost of this gap: "getting parents to come to terms with this new system was initially very difficult…parents were refusing to adopt the new change."

Illustrative Stakeholder Statements on Literacy and Adaptation

Table 1: Stakeholder Statements

Role	Illustrative statement
Parent	"Challenge with the network and the QR scan code…phone could not support the system."
Procurement officer	"Requires me to be on the phone to check for notification most of the time."
Facilitator	"Still used to the traditional way of keeping record."

Administrator	"It took time to transition."
IT administrator	"It was all new to me…extremely difficult."
Management member	"Parents were refusing to adopt the new change."
Administrator	"Things were done manually."
Facilitator	"Sending notes to parents by writing them on paper."

Several stakeholders described falling back on manual, paper-based workarounds during this adjustment period, which is itself a data-handling risk: an administrator noted "things were done manually," a facilitator described "sending notes to parents by writing them on a paper putting it students' bags," and an IT associate described "keeping records of student in paper file" - parallel record-keeping channels operating alongside the digital system precisely where the digital system's controls (unique logins, access logging, encryption expectations) do not apply.

Several stakeholders described falling back on manual, paper-based workarounds during this adjustment period, which is itself a data-handling risk: an administrator noted "things were done manually," a facilitator described "sending notes to parents by writing them on a paper putting it students' bags," and an IT associate described "keeping records of student in paper file" - parallel record-keeping channels operating alongside the digital system precisely where the digital system's controls (unique logins, access logging, encryption expectations) do not apply. A parent recalled the specific device-compatibility trigger for one such workaround: needing to "carry my tablet along to scan when dropping off my kids" because her phone could not support the system, illustrating that literacy and infrastructure gaps are often intertwined rather than cleanly separable - a theme taken up in more depth in a companion analysis of infrastructure constraints [6]. A facilitator described the specific interface complexity underlying much of the reported difficulty: "the system had lots of features, features for teachers were different from that of teachers," while a parent noted a related training gap: "the training was limited and did not cover a large scope so I most of the time took some moments out of my busy schedule to do a personal learning."

Table 2: Summarizes how the coded first-order categories underlying this dataset map onto the two aggregate theoretical dimensions - system usage and conservative users — that structure this paper's analysis.

Aggregate dimension	**First-order codes**
System usage (user benefit)	Verification, communication, data management, assessment, records keeping
Conservative users (user ignorance)	System navigation, mindset, adaptation
Conservative users traditional record keeping)	Physical process, manual record keeping, incompatible device

Coding Structure: First-Order Codes to Aggregate Dimensions

4 DISCUSSION

4.1 The transition window as a security blind spot

These findings support the human-factors framing advanced in Section 3.1. The difficulty was not confined to the least technically engaged stakeholders; it appeared among an IT administrator and a management-level system operator, suggesting that literacy gaps in newly adopted systems are a widespread transitional phenomenon rather than a deficiency limited to specific user groups. More importantly for a security reading of these findings, the parallel-paper-record workaround described by multiple respondents represents exactly the kind of shadow process that undermines a system's formal access controls: information recorded on paper notes carried in a student's bag is not subject to the unique-login, encryption, or third-party-disclosure safeguards that stakeholders elsewhere in this research explicitly valued [4]. In other words, low information literacy did not just slow adoption - during the adoption period, it actively routed sensitive information outside the system's intended security perimeter.

This is consistent with the general information-literacy literature [1,2,3]: institutions that treat literacy-building as a one-off training event, rather than a continuous-adaptation process, risk a security gap precisely during the vulnerable early-adoption window.

4.1.1 Reading DIC through a usability-design lens

Viewed through the usability-design lens introduced in Section 3.1, the DIC findings are consistent with what usability-driven IoT security research would predict: when a system's configuration and everyday-use interface impose a

cognitive burden that exceeds what a non-expert user can readily absorb, the user does not simply fail more slowly - they abandon the intended secure path altogether and substitute an informal one [5]. The QR-code scanning difficulty a parent described, the "lots of features" a facilitator found overwhelming during piloting, and the network-dependent friction a procurement officer experienced are all, in usability terms, configuration and interaction failures rather than motivational ones; none of the affected stakeholders expressed reluctance to use the system securely, only difficulty doing so within the interface and support they were given.

4.1.2 Why this matters more for institutions than for individual consumers

The IoT usability literature this paper draws on was developed primarily with individual consumer devices in mind, where a single user's misconfiguration mainly affects that user's own privacy. The institutional setting examined here is a meaningfully different - and arguably higher-stakes - context: a single stakeholder's literacy gap (a teacher reverting to a paper note, a parent unable to complete digital check-in) does not only affect that individual's own data exposure, it affects the data of the students, and by extension other parents, whose records pass through the same workaround channel. This suggests that institutional information-system usability deserves at least as much dedicated research attention as consumer-device usability has already received, since the security externality of a literacy gap in an institutional setting is borne by third parties (students, other families) who had no part in choosing or configuring the system.

4.1.3 Cross-case comparison with technology-adoption research more broadly

The pattern documented here is broadly consistent with decades of technology-adoption research showing that new-system rollouts produce a temporary competence dip before users reach fluency. What this case study adds is specificity about what fills that dip in a data-sensitive institutional setting: not simply reduced productivity, but a concrete, identifiable substitute practice - paper notes carried in a student's bag, records kept in a paper file - that operates entirely outside the digital system's access controls. This is a more actionable finding for practitioners than a general observation about adoption curves, because it identifies exactly what to look for (parallel paper channels) as an early-warning signal that the literacy gap has become a security gap, rather than leaving practitioners to infer this indirectly from slower task completion or lower system-usage statistics.

4.1.4 Distinguishing literacy risk from resistance to change

It is worth being precise about what this paper's findings do and do not show. None of the difficulty reported by respondents reads as resistance to the underlying goal of digitization; every respondent who described difficulty also described, in the same interview, ways the system had subsequently helped them (better record retrieval, improved communication, reduced administrative burden). The literacy gap documented here is therefore better characterized as a temporary capability gap than as attitudinal resistance, which matters for intervention design: training and support are the appropriate response to a capability gap, whereas a change-management or incentive-based intervention would be the appropriate response to attitudinal resistance, and conflating the two risks applying the wrong remedy.

4.1.5 Literacy gaps, equity, and the digital divide

The literacy gap documented in this case study does not fall evenly across stakeholders, and this unevenness has an equity dimension worth making explicit. Parents who described difficulty tended to attribute it to unfamiliar technology and device limitations rather than to any lack of willingness, and at least one parent's difficulty was compounded by a device-compatibility problem (an older phone that could not run the required application) that is itself a proxy for household resource constraints rather than a pure literacy deficit. This matters because it means literacy interventions designed without attention to the underlying device and connectivity disparities among the parent population risk training people to use a system some of them cannot actually access reliably - a distinction between literacy (knowing how) and access (having the means) that the digital-divide literature has long emphasized in other contexts and that applies with equal force to institutional information-system adoption in a Ghanaian school setting.

4.1.6 Comparative reading against recent human-factors security research

Beyond the IoT usability parallel discussed in Section 3.1, the broader human-factors security literature has repeatedly found that security failures attributed to "user error" are, on closer analysis, more accurately described as design or process failures that made secure behavior unreasonably difficult for the affected user to sustain. The DIC findings are consistent with this reframing: none of the respondents who reverted to paper workarounds expressed a preference for paper over digital record-keeping; all described the digital system, once mastered, as an improvement. The workaround was a response to a temporary capability and interface gap, not a considered preference, which is precisely the pattern human-factors security research identifies as most amenable to correction through better onboarding design rather than through admonishing users to "be more careful."

4.1.7 Toward a general theory of literacy-driven security exposure

Stepping back from the specific findings, this case study points toward a more general theoretical claim worth stating explicitly: in any institutional setting where a digital system replaces a previously manual process, the security posture of that institution during the transition period is bounded not by the digital system's technical controls but by the slower of two rates - the rate at which users acquire sufficient literacy to use the system as intended, and the rate at which the institution eliminates access to the legacy manual alternative. Where the legacy alternative remains easily available, as paper and pen remained available to DIC's staff and parents throughout the PROCARE rollout, the transition period can extend indefinitely for any individual stakeholder who has not yet crossed the literacy threshold, since there is always a lower-effort fallback available. This suggests, counterintuitively, that one lever institutions have not considered in this dataset - deliberately and gradually retiring the manual alternative on a fixed schedule, rather than leaving it available indefinitely as a safety net - might do more to compress the vulnerable transition window than additional training alone, though this would need to be balanced carefully against the risk of leaving genuinely struggling users with no fallback at all during the period before they reach competence.

4.1.8 Interaction with the governance findings

This paper's findings interact directly with the governance findings reported in the companion analysis [4]. The authentication and privacy safeguards that stakeholders described valuing - unique logins, QR-code verification, an expectation of encryption - are only as effective as the consistency with which they are actually used, and this paper's findings show that consistency was measurably compromised during the adoption period specifically by the literacy gap. In other words, the governance gap and the literacy gap are not two independent findings from the same dataset; they compound one another. An institution with strong authentication design but weak onboarding will still leak data through paper workarounds during every future system transition, whether that transition is the initial PROCARE rollout examined here or a subsequent migration to a different or updated system. This compounding relationship is, in the authors' view, the most important single implication connecting this paper to its companion analysis, and it argues for treating governance design and onboarding design as a single, integrated workstream in future institutional technology adoption, rather than as separate IT and training projects run in parallel with limited coordination.

4.2 Future Research Directions

4.2.1 Proposed directions

Building on Section 3.1, future work should test whether AI-assisted, role-specific onboarding - for example, a conversational assistant embedded in a system like PROCARE that answers a parent's or facilitator's question in the moment a difficulty arises, rather than requiring them to recall a one-time training session - reduces both the duration of the adjustment period documented here and the incidence of parallel paper workarounds during it. A second research direction is measuring whether AI-based usage analytics (e.g., identifying which specific stakeholders are repeatedly failing a given system action) can direct scarce training resources to the individuals and roles most at risk of reverting to insecure manual workarounds, rather than applying uniform training to all users regardless of demonstrated need, drawing on techniques already shown to shorten detection times when applied to anomaly monitoring in other operational settings [7]. A third direction, drawn directly from the IoT usability framework discussed above [5], is adapting its proposed evaluation metrics - configuration accuracy rate, error reduction, user decision confidence, adherence to recommended baselines - into a lightweight, low-cost instrument institutions like DIC could administer themselves, without external evaluators, to detect literacy gaps before they manifest as workaround behavior.

4.2.2 Sequencing and feasibility

Of these three directions, the adapted usability-metrics instrument is the most immediately tractable for an institution with DIC's resource profile, since it requires no new technology, only a short structured self-assessment administered during and shortly after onboarding. AI-based usage analytics is more demanding, since, as with the access-log analysis proposed in the companion governance paper [4], it presumes the system already generates structured usage data that most small institutions do not currently collect or retain in a form suitable for analysis. A conversational onboarding assistant is the most resource-intensive of the three, requiring either a vendor-provided feature (most plausibly delivered by the company operating PROCARE itself, which is better resourced than any single client school) or a third-party AI tool integrated into the workflow - making it, like several of the technical directions proposed in the companion papers in this research program, more naturally a vendor responsibility than a single-institution project.

A reasonable pilot sequence, given these feasibility differences, would begin with the self-assessment instrument in a single cohort of newly onboarded institutions, since it produces useful diagnostic data almost immediately and at negligible cost. Results from that pilot - specifically, which stakeholder roles and which specific tasks show the highest error or low-confidence rates - could then directly inform the design of a second-phase usage-analytics study, targeting instrumentation at exactly the tasks the self-assessment identified as highest-risk rather than logging all system activity indiscriminately, which would be both more costly to implement and less immediately actionable. Only once this second

phase had identified specific, recurring failure points would a conversational AI assistant become a well-targeted intervention rather than a generic add-on; building such an assistant before understanding where users actually struggle risks solving the wrong problem, or solving the right problem for the wrong subset of users.

4.2.3 A note on measuring what matters

A risk in adaptive-training research generally is optimizing for a proxy — time-to-competence, session-completion rates — rather than the outcome that actually matters for this paper's argument: reduction in the parallel-workaround behavior that creates security exposure during the transition window. Future evaluation of any AI-assisted literacy tool tested in this setting should therefore include a direct measure of workaround incidence (for example, tracking how many student interactions are still handled via paper note or informal channel at fixed intervals after a new system's introduction) alongside more conventional literacy or satisfaction metrics, since a tool that improves user confidence without measurably reducing workaround behavior would not have addressed the security concern this paper raises.

This point deserves further elaboration because it is easy to overlook in practice. A training program can plausibly succeed on every conventional metric — high attendance, positive post-session feedback, improved quiz scores on system features — while leaving the underlying workaround behavior largely unchanged, if the training addresses declarative knowledge (knowing that a feature exists) rather than the procedural fluency and confidence needed to use that feature under real-world time pressure, with a queue of parents waiting or a classroom of students to manage. The DIC findings are suggestive on this point: several respondents who described initial difficulty also described having received some form of training ("the training was limited and did not cover a large scope"), indicating that training had occurred but had not been sufficient to prevent the workaround behavior this paper documents. This gap between training having occurred and training having been effective is precisely why a direct, outcome-based metric — rather than a training-completion proxy — is necessary for any future evaluation to be credible.

4.3 Practical Implications for Practitioners

Three practical steps follow for institutions in DIC's position. First, onboarding for a new information system should be planned as a graduated process tied to observed proficiency, not a single session delivered once at rollout; the finding that even an IT administrator described the system as "extremely difficult" in the early stages suggests that role seniority or technical background is a poor proxy for actual system familiarity, and training plans should not assume otherwise. Second, institutions should treat the appearance of parallel paper processes during a digital transition as an active signal to intervene, not a harmless bridging habit; each paper note or manually filed record identified in this dataset represents information that briefly or permanently sat outside the digital system's access controls. Third, institutions with a single IT or administrative point of contact — as DIC has — should consider whether literacy support can be distributed (peer support among parents and staff who adapted quickly, for example) rather than relying solely on that one individual, who is otherwise a single point of failure for both technical support and security oversight. Fourth, when selecting a vendor system in the first place, institutions should explicitly evaluate onboarding usability — how much cognitive burden the configuration and everyday-use interface impose on a non-expert user — as a procurement criterion alongside more conventional features and cost, since this case study suggests usability failures translate directly into security-relevant workaround behavior rather than remaining a mere convenience issue.

Fourth, when selecting a vendor system in the first place, institutions should explicitly evaluate onboarding usability — how much cognitive burden the configuration and everyday-use interface impose on a non-expert user — as a procurement criterion alongside more conventional features and cost, since this case study suggests usability failures translate directly into security-relevant workaround behavior rather than remaining a mere convenience issue.

Fifth, and specifically for vendors rather than client institutions: a system such as PROCARE is deployed across many client institutions simultaneously, meaning a vendor-side investment in onboarding usability — clearer in-app guidance, role-specific tutorials, a conversational help assistant — amortizes across every client school rather than requiring each individual institution to solve the same literacy problem independently. This suggests that responsibility for closing the literacy gap documented in this paper should not rest solely with client institutions like DIC; the vendor is better positioned, in terms of both resources and scale economics, to address the root usability issues that generate the workaround behavior this paper documents, echoing the vendor-responsibility argument made in the companion governance and infrastructure papers regarding security tooling and workload protection [4,6].

4.4 Toward a Minimum Viable Onboarding Checklist

Table 3 proposes a minimum viable onboarding checklist for institutions adopting a new information system under resource constraints similar to DIC's. As with the governance checklist proposed in the companion paper [4], each item is chosen to address a specific gap observed directly in the interview data, rather than to replicate a comprehensive enterprise training program that a small institution is unlikely to sustain.

Minimum Viable Onboarding Checklist Derived from DIC Findings

Table 3: Onboarding Checklist

Gap observed	Minimum action
One-time training only	Schedule at least one follow-up session 4–6 weeks after initial rollout
Device incompatibility	Test the system on the actual devices staff and parents will use, before rollout
Paper workaround persists	Track and review any parallel paper records monthly during transition
IT staff assumed competent	Include IT/admin staff explicitly in onboarding, not just end users
No peer support channel	Identify early adopters to informally support slower-adapting peers
No literacy check	Administer a short usability self-assessment at 2 and 8 weeks post-rollout

As with the governance checklist proposed in the companion analysis, this list is intended as an achievable floor rather than a comprehensive training program, on the theory that a small number of consistently applied practices are more likely to be sustained by a resource-constrained institution than a comprehensive but unrealistic training curriculum that is adopted once and then abandoned.

It is worth being explicit about the checklist's sequencing logic. The first two items — a follow-up session and device testing — address root causes that, if left unaddressed, will continue generating new instances of the same workaround behavior regardless of how well the remaining items are executed; they should therefore be treated as prerequisites rather than optional additions. The third item, tracking parallel paper records, functions primarily as a diagnostic: it will not by itself reduce workaround behavior, but it provides the visibility an institution needs to know whether its other interventions are working, in the same spirit as the access-logging prerequisite identified in the companion governance paper's discussion of AI-assisted anomaly detection [4]. The fourth and fifth items address the specific finding, novel to this paper relative to prior literacy research, that technical seniority does not predict system familiarity and that peer-based support may be more available and lower-cost than formal training for a resource-constrained institution. The sixth item, a brief self-assessment, is deliberately lightweight — a full usability audit of the kind conducted in enterprise settings would likely exceed what DIC or a similarly resourced institution could sustain, whereas a short, low-burden check administered twice during the transition period is realistic and still provides an early warning if the transition is taking longer than expected.

A final point on implementation: this checklist, like the governance checklist in the companion paper, is designed to be owned collectively rather than assigned to a single individual, precisely because the findings in Section 3.2 show that literacy gaps affected even the single IT staff member who might otherwise be assumed to own an onboarding checklist by default. Distributing ownership — a management member responsible for the follow-up session, an IT associate responsible for device testing, a facilitator responsible for tracking paper records in their own classroom — reduces the single-point-of-failure risk this paper's findings suggest is otherwise inherent in leaving onboarding entirely to one role.

4.5 Designing a Workaround-Incidence Metric

Section 3.4 argued that future evaluation of literacy interventions should measure workaround incidence directly rather than relying on proxy measures such as training completion or self-reported confidence. This section sketches, in more concrete terms, what such a metric might look like for an institution in DIC's position, since the absence of an existing standard instrument for this purpose is itself a gap this paper's argument implies needs filling.

A minimal workaround-incidence metric could be constructed from three simple, low-burden observations, each collected at fixed intervals (for example, weeks 2, 6, and 12 following a new system's introduction). First, a count of paper-based records created for any function the digital system is intended to handle - attendance notes, pick-up authorizations, behavioral feedback - gathered simply by asking each facilitator or teacher to retain and tally any such paper records they produce during a defined observation week, rather than through any technical instrumentation. Second, a short structured question posed to each stakeholder role, asking them to identify which specific system function, if any, they still find themselves avoiding or completing manually, mirroring the open-ended interview approach used in the underlying dataset but administered as a brief recurring check-in rather than a one-time interview. Third, where available, a simple comparison between the number of active user accounts and the number of accounts that have completed at least one full transaction cycle (for example, a parent account that has completed at least one digital pick-up authorization), which would surface accounts that exist nominally but have not yet transitioned to actual digital use.

None of these three measures requires specialized technical infrastructure, which is deliberate: a metric that itself requires the same technical sophistication the institution is still building toward would not be usable during the exact period it is meant to monitor. A composite score combining all three - a rough, non-technical "digital transition index" - could then be tracked over successive weeks to establish an empirical timeline for when an institution's workaround behavior falls to a stable, low baseline, providing both a practical management tool and, potentially, a standardized instrument other researchers could adopt to make workaround-incidence findings comparable across institutions and settings in future research, addressing the specific gap in comparability that Section 3.7 identifies as a limitation of this single-institution study.

Limitations

This paper's findings are drawn from stakeholder self-reports of difficulty and adaptation during a single system rollout at one institution; the specific incidence of the "parallel paper workaround" pattern was not independently measured (for instance, by auditing how many student records existed only on paper at a given point in time), so its scale relative to properly digitized records cannot be quantified from this dataset. The information-literacy citations underpinning Section 3.1 (Thanuskodi, 2019; Corrall, 2008; Koltay, 2017) are reproduced from the original case-study report's in-text citations but, as noted in the reference list below, full bibliographic entries for these could not be independently verified during this revision; readers seeking to build on this literature review should independently locate and confirm full source details before relying on them for further citation. The usability-metrics framework drawn from IoT security research [5] was developed for consumer device configuration rather than institutional information systems, and while this paper argues the underlying usability principles transfer, the specific metrics proposed have not been tested in an institutional setting and would require adaptation and validation before use.

A further limitation concerns generalizability beyond the education sector. DIC's stakeholder mix - parents balancing the system against child-safety concerns, teachers and facilitators with varying technical backgrounds, a small administrative core - is specific to a school setting, and institutions in other sectors (healthcare, small business, community organizations) may exhibit different literacy-risk patterns depending on their own stakeholder composition, regulatory environment, and the specific system being adopted. The core mechanism this paper identifies - that literacy gaps manifest as identifiable workaround behaviors during a transition window, and that these workarounds carry security consequences - is plausibly generalizable in structure, but the specific workaround observed here (paper notes, manual filing) is particular to a setting where a paper-based predecessor process existed and remained institutionally familiar; a setting without such a predecessor process might exhibit a different, though still security-relevant, substitute behavior.

A final limitation concerns the workaround-incidence metric proposed in Section 3.7: it is a conceptual proposal derived from this paper's findings, not an instrument that has been piloted, validated, or tested for reliability across observers or institutions. Before any institution adopts it as a management tool, and before any researcher adopts it as a basis for cross-institutional comparison, it would need the kind of validation work - inter-rater reliability checks, testing against an independently verified count of actual paper records, piloting across more than one institutional setting - that this single-institution qualitative study was not designed to provide.

5 CONCLUSION

Divine International College's transition to PROCARE shows that information literacy operates as a security variable, not merely an efficiency one: the period during which stakeholders - including technically responsible staff - struggled with the system was also the period in which sensitive information was most likely to be handled outside the system's intended safeguards, via paper workarounds. Institutions introducing new information systems should budget for structured, role-specific, and ideally adaptive onboarding as a security investment on par with technical controls, not a separate training-and-development line item. Future AI-assisted literacy tools, informed by usability-driven design principles already developed for consumer IoT security [5], may make this affordable for resource-constrained institutions, but, as with the compliance tooling discussed in the companion governance paper [4], this remains a research proposition to be tested rather than a proven solution.

For institutional leaders, the single most actionable recommendation from this study is also the simplest: treat the weeks immediately following any new system's introduction as a defined, monitored security period, not merely an operational adjustment period. During that window, actively look for the specific signature this paper identifies - paper notes, manual filing, informal workarounds - rather than waiting for a conventional security incident to reveal that the digital system's controls were being bypassed. For researchers, this study's main contribution is the reframing itself: literacy and security, in the specific context of new-system adoption, are not separate research questions answered by separate literatures, but a single phenomenon that should be studied and measured together, using instruments like the workaround-incidence metric proposed in Section 3.7 as a starting point for future, more rigorously validated work. For vendors, finally, the implication is that onboarding usability is not a peripheral product-design concern but a core

security feature of any information system marketed to resource-constrained institutions, deserving the same design attention and evaluation rigor currently reserved for authentication and encryption features. A vendor that treats onboarding as an afterthought is, in effect, exporting a security liability to every client institution it serves, since each institution must then independently absorb the literacy-gap risk this paper documents, rather than benefiting from a solution engineered once, well, and at scale by the party best positioned to engineer it. The central claim of this paper is narrower than a general call for more training: it is that literacy gaps have a specific, identifiable security signature - the parallel paper workaround - and that future interventions, whether human-delivered or AI-assisted, should be evaluated against their ability to reduce that specific signature, not merely against general measures of user satisfaction or confidence.

More broadly, this case study suggests that the conventional separation between "training" and "security" budgets in institutional planning is analytically unsound, at least for the transition period following a new system's adoption. If a literacy gap reliably produces a security-relevant workaround, as this dataset suggests, then onboarding investment is not merely a productivity expenditure that happens to have some indirect security benefit - it is a direct security control, comparable in function, if not in mechanism, to the authentication and access-control measures examined in the companion governance paper [4]. Institutions, vendors, and researchers who continue to treat these as separate budget lines and separate research literatures risk underinvesting in exactly the intervention - usable, well-sequenced onboarding - that this case study suggests would do the most to close the specific security gap documented here.

Finally, this paper's findings should be read as a call to extend, rather than replace, the existing information-literacy and human-factors security literatures. What this case study contributes is not a new theory of literacy or a new theory of security in isolation, but an empirical demonstration of exactly how the two intersect in a single, resource-constrained institutional setting: literacy gaps do not merely slow adoption, they open a specific, identifiable, and - this paper argues - measurable channel through which an institution's data-handling practice diverges from its intended, governed practice. Closing that channel is not solely a training problem or solely a security problem; it is both at once, and future research, institutional practice, and vendor product design would all benefit from treating it that way.

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] S. Thanuskodi, "Information literacy competency and data management practices among information professionals," as cited in the source case-study report, 2019.

[2] S. Corrall, "Information literacy strategy development in higher education," as cited in the source case-study report, 2008.

[3] T. Koltay, "Information literacy and data management practices," as cited in the source case-study report, 2017.

[4] Independent Research Study, "Balancing access and privacy: Regulatory gaps, authentication, and data security in the digital preservation of institutional records," companion paper derived from the same case-study dataset, 2026.

[5] Independent Research Study, "Beyond bandwidth: Infrastructure constraints, cost barriers, and the case for AI-assisted secure digital preservation in resource-limited institutions," companion paper derived from the same case-study dataset, 2026.

[6] K. Boamah, "Usability standards for privacy-preserving security configuration in IoT devices," *International Journal of Research Publication and Reviews*, vol. 6, no. 11, 2025. doi: 10.55248/gengpi.06.1125.3973.

[7] K. G. Boamah, "Cloud Security Posture Management: A comprehensive analysis of automated risk identification and mitigation in multi-cloud environments," *International Journal of Research and Innovation in Social Science*, vol. 9, no. 11, pp. 4458–4471, 2025. doi: 10.47772/IJRISS.2025.91100349.

[8] A. Asante, "Strengthening healthcare cloud security using Cloud Workload Protection Platforms (CWPP): A framework for protecting patient-critical workloads in health data warehouses," *International Journal of Research and Innovation in Social Science*, vol. 9, no. 11, pp. 7863–7889, 2025. doi: 10.47772/IJRISS.2025.91100613.

[9] A. Asante, "Evaluating confidential computing runtimes to enforce verifiable privacy guarantees and automated GRC controls in multi-tenant cloud data processing workflows," *International Journal of Computer Applications Technology and Research*, vol. 14, no. 11, pp. 40–48, 2025. doi: 10.7753/IJCATR1411.1005.

[10] K. G. Boamah, A. Asante, A. Timean, and K. F. Okai, "Artificial intelligence integration in cyber incident response teams to enable faster containment, forensic accuracy, and resilient business continuity," *International Journal of Science and Research Archive*, vol. 17, no. 1, pp. 1263–1280, 2025. doi: 10.30574/ijsra.2025.17.1.2933.