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

BEYOND BANDWIDTH: INFRASTRUCTURE CONSTRAINTS, COST BARRIERS, AND AI-ASSISTED SECURE DIGITAL PRESERVATION IN RESOURCE-LIMITED INSTITUTIONS

Afua Asantewaa Asante *

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

ORCID Details
Afua Asantewaa Asante: https://orcid.org/0009-0001-5634-5759

International Journal of Science and Research Archive, 2026, 20(02), 335–346

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.2026.20.2.1560

ABSTRACT

Digital preservation research has traditionally addressed infrastructure and cost as adoption barriers to be overcome on the way to a stable end state. This study argues, using qualitative case-study evidence from a Ghanaian school information system, that infrastructure limitations are not merely a one-time adoption hurdle but an ongoing operational and security condition that resource-constrained institutions manage continuously. Stakeholders across roles described system slowdowns tied to unreliable internet connectivity, the cost of acquiring and maintaining compatible devices, and the need for regular system maintenance and backups, all recurring themes rather than one-off implementation issues. The findings are situated within literature on internet connectivity, digitisation cost, and long-term data accessibility, and it is argued that infrastructure constraints have a direct security dimension: institutions that cannot reliably back up, patch, or monitor their systems are more exposed to data loss and unauthorised access regardless of how well-designed their access controls are on paper. The study closes with a discussion of artificial-intelligence-assisted, low-bandwidth approaches, including lightweight anomaly detection, automated backup scheduling, predictive maintenance, and cloud workload protection adapted for intermittent-connectivity settings, as a promising but unproven direction for helping infrastructure-constrained institutions sustain secure digital preservation without proportionally large capital investment.

Keywords: Digital Preservation; Infrastructure; Internet Connectivity; Data Security; Artificial Intelligence

1 INTRODUCTION

Discussions of digital preservation adoption often treat infrastructure and cost as a threshold to cross: once an institution has acquired adequate connectivity and hardware, the assumption goes, digitization proceeds and the infrastructure question recedes. Case-study evidence from Divine International College (DIC), a Ghanaian school using the PROCARE information management system, suggests otherwise. Infrastructure limitations - connectivity, device compatibility, cost of maintenance - recurred as live, ongoing concerns well after initial adoption, not as issues resolved by the original purchase decision. This paper argues that in resource-constrained institutional settings, infrastructure

adequacy is better understood as a continuous operating condition with direct security consequences, rather than a one-time capital expenditure.

This paper makes four contributions. First, it provides empirical evidence that infrastructure and cost constraints in a small-institution setting recur throughout a system's operational life rather than concentrating at the point of adoption, contrary to how digitization cost is often modeled in prior literature. Second, it connects infrastructure adequacy directly to security posture, arguing that an institution's practical capacity to patch, back up, and monitor its systems - which this case shows to be connectivity- and budget-constrained - is itself a security variable, not merely an operational convenience. Third, it introduces a total-cost-of-ownership framing to digitization cost research, showing that unbudgeted compatibility purchases and ongoing maintenance burden are cost categories distinct from, and in this case study at least as significant as, initial procurement cost. Fourth, it proposes a sequenced research agenda for AI-assisted, low-bandwidth monitoring and maintenance tools, explicitly designed for the connectivity and staffing constraints documented in this case.

This paper revisits the infrastructure and cost dimension of a broader case study of digital preservation and data management at DIC, reframing it alongside companion analyses of governance/authentication [1] and information literacy [2] produced from the same dataset, and extends the discussion toward AI-assisted approaches suited to low-bandwidth, cost-constrained environments. It is worth being explicit about scope at the outset, as in the companion papers: this paper does not claim that DIC's infrastructure planning was negligent or that its stakeholders failed to anticipate reasonable problems. The evidence, if anything, shows the opposite - when device incompatibility emerged as a real operational problem, the institution responded, purchasing additional compatible hardware for key staff. What this paper argues is narrower: that even a reasonably responsive institution, operating without a total-cost-of-ownership framework or a documented infrastructure resilience process, will find itself perpetually reacting to infrastructure problems as they arise rather than anticipating and budgeting for them in advance, and that this reactive posture has security consequences beyond the immediate operational inconvenience it is usually understood to represent. The remainder of the paper proceeds as follows. Section 3.1 reviews literature on connectivity, digitization cost, long-term accessibility, and AI for constrained environments. Section 2 describes the 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 infrastructure checklist, Section 3.7 discusses limitations, and Section 3.8 concludes.

2 MATERIAL AND METHODS

2.1 Research paradigm and case

This paper draws on the same interpretivist and constructivist qualitative paradigm, case institution, and sampling strategy described in the companion papers [1,2]: 25 of 27 targeted stakeholders interviewed (83.3%) at Divine International College, selected via a maximal-variation, purposive and convenient sampling strategy, coded using a first-order/second-order/aggregate-dimension approach. This paper isolates two aggregate dimensions relevant to infrastructure and cost - "system performance" (system jams, internet connectivity, implementation cost, user interface, system updates) and "data management/accessibility" (data retrieval, tracking, access, preservation) - for focused analysis.

2.2 Reliability, validity, and reflexivity

The same reliability and validity procedures described in the companion papers apply here: 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 in the companion analyses, the researchers' existing relationship with the case institution likely eased candor about infrastructure difficulties and cost constraints that an external, unfamiliar research team might not have elicited as readily, while also meaning the coding categories applied reflect the research team's own analytic framing.

2.3 Reflexivity and interview protocol

As in the companion papers [1,2], the interview protocol underlying this dataset used open-ended rather than checklist-style questions, asking respondents to describe their experience of the system's performance and cost in their own words. This is likely why infrastructure and cost complaints in this dataset are specific and situational ("the system jams up and it needs strong internet connection") rather than generic ratings of satisfaction, which would have obscured the operational detail this paper relies on. The researchers' existing relationship with DIC, noted in the companion papers, likely eased candor about cost overruns and performance problems that an institution might otherwise be reluctant to disclose to an unfamiliar outside researcher, while also meaning the coding categories applied ("system performance," "data management/accessibility") reflect the research team's own analytic framing rather than categories respondents

would necessarily have generated unprompted. Member-checking, described above, was used specifically to guard against this risk.

3 RESULTS AND DISCUSSION

3.1 Literature Context

3.1.1 Connectivity as a constraint on digital preservation

Prior research on internet connectivity and digital content management finds that slow or unstable internet connections make it harder to upload, download, or stream digital content, directly diminishing the efficacy of preservation efforts [3]. Work on privacy-preserving frameworks for electronic records likewise notes that connectivity and infrastructure gaps can delay the transmission of data to preservation systems or impede retrieval when information is needed [4]. Early and still-cited treatments of digital libraries emphasize that inadequate internet connectivity constrains not just preservation itself but users' practical ability to access and use preserved content [5] - meaning infrastructure gaps degrade the value of digitization even when the underlying preservation strategy is otherwise sound.

3.1.2 The cost of digitization and maintenance

Comparative studies of large-scale digitization projects [6] and work on records management governance in Africa [7] both point to the high cost of acquiring and maintaining digitization equipment as a constraint on the scale and pace of digitization efforts, particularly for institutions operating under tight budgets. This cost pressure does not end at the point of adoption: ongoing maintenance, software updates, and device replacement represent a recurring cost stream that resource-constrained institutions must sustain indefinitely to keep a digitized system secure and functional.

3.1.3 Long-term accessibility and data management outcomes

Where infrastructure and cost constraints are managed successfully, the literature documents clear benefits. Work on digital preservation and access argues that digital preservation maintains the accessibility and availability of important information resources over the long term [8], while more recent work on academic library preservation notes that adopting the right techniques and technologies protects digital material from technical obsolescence, data deterioration, and loss [9] - allowing institutions to avoid losing information to format changes, device obsolescence, or medium degradation. These benefits, however, presuppose an institution's ongoing ability to fund and maintain the infrastructure that delivers them, which is precisely the condition this paper's case-study evidence calls into question for smaller, resource-constrained institutions.

3.1.4 Infrastructure limitations as a security condition

A connection largely unexamined in the original case study, but significant for a security-oriented reading of these findings, is that infrastructure limitations are also security limitations. An institution that cannot reliably back up its data, apply timely software patches and updates, or monitor system performance because of connectivity or budget constraints is more exposed to data loss, corruption, and, potentially, unauthorized access, independent of how well-designed its authentication and access-control policies are on paper (see the companion governance analysis [1]). Security controls that depend on regular maintenance and monitoring are only as strong as an institution's practical ability to perform that maintenance consistently.

3.1.5 Total cost of ownership versus procurement cost

Much of the digitization-cost literature, including comparative project studies [6] and governance-focused work [7], focuses on the cost of acquiring digitizing equipment. The DIC findings suggest a total-cost-of-ownership framing is more appropriate for resource-constrained institutions: procurement cost, unbudgeted compatibility purchases, ongoing connectivity cost, and staff time spent on manual workarounds during outages or slowdowns are all cost categories distinct from the initial licensing or hardware spend, yet all were raised, in one form or another, by DIC respondents. A total-cost-of-ownership framing also clarifies why AI-assisted approaches merit research attention: their value proposition is not necessarily a lower headline price, but a potential reduction in the ongoing, harder-to-budget cost categories that this case study suggests are where resource-constrained institutions are actually strained.

3.1.6 Cloud security posture and workload protection for constrained environments

Recent work on cloud security posture management (CSPM) shows that automated resource discovery, continuous compliance monitoring, and automated remediation can reduce misconfiguration incidents by 60–80% and cut threat-detection time from weeks to hours relative to manual oversight, evaluated across enterprise multi-cloud deployments spanning AWS, Azure, and GCP [10]. Related work on Cloud Workload Protection Platforms (CWPP) for healthcare data warehouses 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]. Neither line of work was designed with intermittent-connectivity, single-IT-staff institutions like DIC in mind; both presume a level of infrastructure and technical capacity well beyond what this case study documents. This gap — between what enterprise-grade security tooling assumes and what resource-constrained institutions can actually sustain — is a central motivation for the AI-assisted, low-bandwidth research agenda proposed in Section 3.4.

3.1.7 Artificial intelligence for low-bandwidth, cost-constrained environments

A promising but underexplored direction for institutions facing the constraints documented here is the use of AI-assisted tools specifically designed for low-bandwidth, low-budget environments, rather than the well-resourced enterprise environments most preservation and security tooling is built for. Possibilities include lightweight, on-device anomaly detection that does not require constant cloud connectivity to flag irregular system behavior; predictive-maintenance models that use available usage data to anticipate when a system is likely to fail or require intervention, allowing an institution with a single IT associate (as at DIC) to prioritize limited attention; and automated backup scheduling that adapts to intermittent connectivity, backing up opportunistically when bandwidth is available rather than requiring a constant connection. Recent work applying AI to cyber incident response teams shows that AI-driven analytics can continuously monitor system behavior, detect subtle anomalies, and rapidly correlate multi-source indicators of compromise even in complex, fast-evolving environments [12], suggesting the underlying analytic techniques are mature enough to be adapted downward to smaller-scale, lower-connectivity deployments, even though the specific tooling examined in that work was designed for well-resourced incident-response teams rather than a single school IT associate. None of these approaches has been tested in this case study's setting, and their practical viability — particularly cost, and the technical capacity required to deploy and maintain them — is an open question rather than a settled recommendation.

3.1.8 Infrastructure constraints in the Ghanaian and regional context

The connectivity and cost constraints documented in this case study should be read against the broader infrastructure context of Ghana and comparable Sub-Saharan African settings, where internet connectivity quality and cost vary substantially between urban and peri-urban areas and where import costs for compatible devices can represent a significant share of an institution's technology budget relative to local wage levels. DIC, situated in Accra, is arguably better positioned on connectivity than many rural Ghanaian institutions would be, which suggests that the connectivity-related findings reported here may understate the severity of the problem for less urban institutions attempting similar digitization efforts. This regional context also explains why device-compatibility problems recurred as forcefully as they did in this dataset: a diverse population of parents and staff, each using whatever device they already own rather than a standardized institution-issued device, is a more cost-realistic deployment model for a Ghanaian private school than the assumption of standardized hardware common in enterprise information-system design, and this mismatch between deployment assumption and actual device diversity is itself a source of the compatibility friction respondents described.

Findings

Connectivity problems recurred across multiple respondents and roles. A facilitator stated plainly: "the system jams up and it needs strong internet connection." A visual-arts facilitator echoed this: "the system jams and the scanning were also an issue." A teacher, asked about system flexibility, pointed simply to "internet use" as the defining factor. These are not isolated complaints but a structural characteristic of operating a cloud-connected system in an environment where connectivity cannot be assumed.

Cost surfaced both as an implementation and an ongoing burden, summarized alongside connectivity findings in Table 1. A management member described a targeted capital response: "the school ended up buying some mini iPads for key users like the administration office, the IT department, the library and administration" - additional, unbudgeted hardware spend necessitated by device incompatibility. A procurement officer identified "cost of implementing and maintaining" the system as a distinct concern from the purchase price alone, and separately noted that the system "requires ongoing maintenance and updates to ensure that the system remains up-to-date" - explicitly framing maintenance, not just acquisition, as a resource commitment.

3.2 Illustrative Stakeholder Statements on Infrastructure and Cost

Table 1: Stakeholder Statements

Role	Illustrative statement
Facilitator	"The system jams up and it needs strong internet connection."
Teacher	"Internet use" (as the defining factor in system flexibility).
Management member	"The school ended up buying some mini iPads for key users."
Procurement officer	"Cost of implementing and maintaining" as distinct from purchase price.
Parent	"Carry my tablet along to scan when dropping off my kids."
Teacher/Facilitator	"Regular updates" / "Automatic software update."
IT associate	"Track list of students…over a period of time."
Facilitator	"It us helped us not lose any of our records."

Device incompatibility compounded the cost and connectivity issues: a parent recalled needing to "carry my tablet along to scan when dropping off my kids" because her phone could not support the system, while another parent noted more optimistically that "the system works on both android, windows and iOS" — though this reassurance sat alongside multiple other respondents' reports of practical incompatibility. On the maintenance side, a facilitator and a teacher both raised "regular updates" and "automatic software update" as recurring needs, and a procurement officer specifically linked ongoing maintenance to system currency: the system "requires ongoing maintenance and updates to ensure that the system remains up-to-date."

Despite these constraints, respondents were clear that the system delivered substantial data-management benefits once operating: a teacher noted "there is increased accessibility and improved security," an IT associate observed that the school could now "track list of students who reported in a particular year or grade and number of withdrawn students over a period of time," and a facilitator credited the system with helping the institution avoid losing records entirely: "it us helped us not lose any of our records."

Table 2: Summarizes how the coded first-order categories underlying this dataset map onto the two aggregate theoretical dimensions - system performance and data management/accessibility - that structure this paper's analysis.

Aggregate dimension	First-order codes
System performance	System jams, internet connectivity, implementation cost, user interface, system updates
System performance (maintenance)	Limited capacity for digitization, system maintenance, system integration
Data management/accessibility	Data retrieval, data tracking, data access

Coding Structure: First-Order Codes to Aggregate Dimensions

A facilitator's account of the piloting stage illustrates how performance and cost concerns compounded one another in practice: "for the piloting stage it was difficult because I was not able to navigate my way around the system. The system had lots of features, features for teachers were different from that of teachers." A parent's account of an early connectivity failure similarly shows the interaction between infrastructure and everyday use: difficulties with "the network and the QR scan code" were compounded by a device that "could not support the system," requiring an ad hoc hardware workaround rather than a software fix.

4 DISCUSSION

4.1 Supporting and extending prior literature

These findings support prior work on connectivity's impact on preservation-system performance [3,4] and on the cost burden of digitization [6,7] - but they extend that literature by showing these constraints as continuous operating conditions at DIC, not resolved adoption hurdles. The additional, unbudgeted iPad purchase described by one

respondent illustrates that infrastructure cost is elastic: it expands as new incompatibilities surface, rather than being fixed at the point of initial procurement.

Read through the security lens introduced in Section 3.1, the maintenance and update burden repeatedly raised by respondents is not simply an IT inconvenience - a system that is not regularly patched or backed up because a single IT associate lacks the bandwidth or connectivity to do so consistently is a system with a growing, unaddressed security exposure, regardless of how sound its original authentication design was (see the companion governance paper [1]). The long-term accessibility benefits documented in prior literature [8,9] depend on an institution's sustained capacity to perform exactly this kind of maintenance - a capacity this case study suggests cannot be assumed for smaller, resource-constrained institutions.

4.2 Comparative reading against enterprise cloud security tooling

Placing DIC's findings alongside the CSPM and CWPP literature reviewed in Section 3.1 highlights a significant capability gap rather than a difference in underlying need. Enterprise CSPM tooling achieves its 60–80% reduction in misconfiguration incidents [10] through automated, continuous monitoring that presumes reliable connectivity and a technical team capable of interpreting and acting on its output. DIC has neither: connectivity is explicitly described by respondents as intermittent, and the institution's entire technical capacity rests with a single IT associate who also handles day-to-day system operations. This is not evidence that DIC does not need the kind of continuous visibility CSPM and CWPP tooling provide - if anything, the recurring, unaddressed maintenance burden described by respondents suggests DIC needs it more, precisely because it lacks the staffing to compensate through manual vigilance the way a larger institution might. The gap is one of tooling designed for the wrong resource profile, not one of need.

This gap is worth dwelling on further because it illustrates a more general mismatch in how security and infrastructure tooling markets tend to develop. Commercial CSPM and CWPP products are built, tested, and priced for organizations that can afford both the tooling and the staff to operate it - typically enterprises with dedicated security teams and predictable, high-bandwidth connectivity. The evaluation metrics used in that literature (reduction in misconfiguration incidents, threat-detection time measured in hours) implicitly assume a baseline of technical capacity and connectivity that simply does not exist at an institution like DIC. This is not a criticism of that literature on its own terms - it is measuring real, valuable improvements for the settings it was designed to serve - but it does mean that resource-constrained institutions cannot simply wait for enterprise tooling to become cheap enough to adopt as-is; the tooling itself needs to be redesigned around a different set of baseline assumptions, which is precisely the research agenda proposed in Section 3.4.

4.3 The elasticity of infrastructure cost

A theme that deserves separate emphasis is the elasticity of infrastructure cost demonstrated in this case: the mini-iPad purchase was not part of any original procurement plan but was triggered reactively once device incompatibility became an operational problem. This pattern - cost expanding in response to newly discovered incompatibilities rather than being fixed at procurement - has an important implication for how institutions and researchers model digitization cost. A total-cost-of-ownership estimate produced at the point of system selection, without a contingency allowance for this kind of reactive expansion, will systematically underestimate the true cost resource-constrained institutions will bear over a system's operational life.

4.4 Implications for procurement practice

This elasticity finding has a direct, actionable implication for how institutions like DIC should approach procurement decisions in the first place. Rather than evaluating a candidate information system purely on its advertised licensing or subscription cost, procurement decision-makers should explicitly budget a compatibility contingency - informed by a pre-purchase device-compatibility audit of the actual devices staff and parents currently own, rather than an assumption of standardized, institution-issued hardware. DIC's experience suggests this audit would have identified the device-incompatibility problem before it became a reactive, unbudgeted expense; a parent's phone being unable to run the required scanning application, or a staff member's need to carry a separate tablet, are precisely the kind of concrete, discoverable facts a modest pre-purchase survey of device ownership among the parent and staff population could have surfaced in advance. This is a low-cost procurement step - essentially a short survey - relative to the reactive cost it could help an institution avoid, and it follows directly from treating total cost of ownership, rather than sticker price, as the correct basis for procurement decisions in resource-constrained settings.

4.5 Cross-case comparison with prior digitization-cost research

The pattern documented here is broadly consistent with, and adds granularity to, prior digitization-cost literature. Lapworth et al.'s [6] comparison of three large-scale digitization projects and Mosweu and Rakemane's [7] governance-focused analysis both identify cost as a constraint on digitization scale and pace, generally framed at the level of institutional or national budgeting decisions. DIC's data show what that constraint looks like from inside a single small

institution's day-to-day operational reality: cost pressure does not arrive as a single, anticipated budget line, but as a series of unplanned, reactive expenditures - the mini-iPad purchase being the clearest example - each triggered by a specific, previously unforeseen incompatibility or performance failure. This granular, reactive pattern is arguably more representative of how cost actually functions as a barrier for small institutions than the more static, planning-stage cost estimates typically discussed in prior digitization-cost literature.

4.6 Distinguishing infrastructure risk from adoption resistance

As with the literacy findings discussed in a companion paper [2], it is worth being precise about what the infrastructure and cost findings do and do not show. No respondent expressed a preference for reverting to manual, non-digital operations because of infrastructure difficulty; every respondent who reported a connectivity or cost problem also described, elsewhere in the same interview, benefits the system had delivered once operational. This distinguishes infrastructure risk, as documented here, from any suggestion that DIC's stakeholders were reluctant adopters - the constraints reported are genuine operational limitations on an otherwise valued system, not evidence of resistance to digitization as a goal.

4.7 Interaction with the governance and literacy findings

This paper's findings interact directly with those reported in the two companion analyses [1,2]. The single IT associate responsible for maintenance and updates at DIC is the same role identified in the governance paper as responsible for reviewing user access and in the literacy paper as itself experiencing significant onboarding difficulty during the PROCARE rollout. This concentration of responsibility in one role, already flagged as a governance risk and a literacy risk in the companion papers, is compounded further by this paper's infrastructure findings: the same individual who must review access logs and who struggled with the system's own learning curve is also the person responsible for applying updates, managing backups, and responding to connectivity-related performance problems. A single point of failure that appears three times across three independent analytic lenses - governance, literacy, and infrastructure - is a stronger and more urgent finding than any one lens would suggest in isolation, and it is, in the authors' view, the single most important cross-cutting theme to emerge from this three-paper research program.

4.8 Future Research Directions

4.8.1 Proposed directions

Building on Section 3.1, three research directions are proposed for institutions facing DIC's combination of intermittent connectivity, limited IT staffing, and constrained budgets. First, predictive-maintenance research should test whether lightweight, locally-run models can flag system components likely to fail or require attention, allowing a single IT associate to prioritize effort rather than respond only after a system jam already disrupts operations. Second, backup and synchronization research should test opportunistic, connectivity-aware backup scheduling that performs backups when bandwidth allows, rather than requiring constant connectivity - directly addressing the "system jam …needs strong internet" pattern reported here. Third, cost-effectiveness research should directly compare the total cost of AI-assisted monitoring and maintenance tooling against the ad hoc capital responses documented in this case study (such as the unbudgeted iPad purchase), to establish whether AI-assisted approaches genuinely reduce total infrastructure cost for small institutions or merely relocate it. None of these directions should be read as an endorsement of any specific product or vendor; they are proposed as an empirical research agenda.

4.8.2 Adapting enterprise tooling downward

A fourth, cross-cutting direction follows directly from the comparative reading in Section 3.3: rather than building entirely new tools from scratch for resource-constrained settings, researchers should investigate whether the underlying analytic techniques behind CSPM [10], CWPP [11], and AI-assisted incident response [12] can be adapted downward - simplified, made to tolerate intermittent connectivity, and designed to produce actionable output for a single generalist IT associate rather than a specialized security team. This is a more tractable near-term goal than expecting resource-constrained institutions to adopt enterprise tooling as built, and it treats the existing enterprise research as a starting point to adapt rather than a benchmark resource-constrained institutions are expected to somehow reach unaided.

Concretely, this downward adaptation would likely involve at least three simplifications relative to the enterprise originals. First, replacing continuous, always-on monitoring with periodic, connectivity-aware checks that run opportunistically whenever the system detects an available connection, rather than assuming constant connectivity as CSPM tooling typically does. Second, replacing dashboards and alerts designed for a specialized security analyst with plain-language summaries and a small, fixed set of recommended actions appropriate for a generalist IT associate without formal security training - echoing the usability-first design philosophy already argued for in the companion literacy paper's discussion of onboarding [2]. Third, replacing the workload-level granularity of enterprise CWPP tooling, which distinguishes between many discrete containerized workloads, with a simpler, whole-system view

appropriate for an institution running a single vendor-supplied information system rather than a complex multi-service cloud architecture. None of these simplifications are technically exotic; the research contribution would lie less in novel algorithms and more in the deliberate, resource-aware redesign of existing, already-proven techniques for a different operating context.

4.8.3 Sequencing and feasibility

Of these four directions, opportunistic backup scheduling is the most immediately tractable, since it requires no new analytic capability, only a scheduling and connectivity-awareness layer added to existing backup routines. Predictive maintenance is more demanding, requiring sufficient historical usage or performance data to train even a lightweight model - data that DIC, like many small institutions, may not currently be capturing in a usable form. Downward-adapted CSPM/CWPP tooling is the most ambitious of the four, since it presumes not only technical adaptation but a business case for a vendor or research team to invest in serving a market (small, low-connectivity institutions) that is less lucrative than the enterprise market such tooling currently targets - making this, like several directions proposed in the companion papers, more a matter of vendor incentive alignment than a purely technical research problem.

A reasonable pilot sequence would begin with opportunistic backup scheduling deployed at a small number of volunteer institutions similar to DIC, since it produces an immediately measurable outcome (backup success rate under real-world intermittent connectivity) with minimal implementation risk. Results from that pilot - specifically, how much connectivity variability the scheduling approach can tolerate before backup success rates degrade meaningfully - would then usefully inform the design specification for predictive-maintenance research, since the same connectivity-variability data collected during the backup pilot could double as the historical usage data a predictive-maintenance model would need for training. Only after both of these more tractable directions had been piloted and shown to work reliably in a genuinely low-connectivity setting would downward-adapted CSPM/CWPP tooling become a reasonable next investment, since building sophisticated posture-management tooling before establishing that more basic connectivity-aware infrastructure functions reliably would risk layering complexity on top of an unstable foundation.

4.8.4 A note on measuring what matters

As with the workaround-incidence metric proposed for the literacy dimension of this research program, a risk in infrastructure research generally is optimizing for a proxy - uptime percentage, number of tickets resolved - rather than the outcome that actually matters for this paper's argument: whether an institution's practical capacity to back up, patch, and monitor its systems is keeping pace with the sensitivity of the data those systems hold. A system can show excellent uptime statistics while still carrying substantial security exposure if, for example, backups are running successfully but have never been tested for successful restoration, or if software updates are being applied inconsistently because the single available IT associate is stretched across too many competing responsibilities. Future evaluation of any AI-assisted infrastructure tool proposed in Section 3.4 should therefore be judged against its ability to close this specific gap between surface-level operational metrics and underlying security-relevant capacity, rather than against uptime or ticket-resolution statistics alone, which is precisely the distinction the infrastructure resilience metric proposed in Section 3.5 is designed to capture.

4.9 Practical Implications for Practitioners

Three implications follow for institutions facing similar constraints. First, device-compatibility contingency should be budgeted explicitly rather than discovered reactively; DIC's unbudgeted iPad purchase for key staff illustrates that hardware cost does not end with a system's initial licensing or subscription fee, and institutions should build a compatibility-testing step into any system-selection process before rollout, not after. Second, because a single IT associate appears to carry responsibility for maintenance, updates, and backups at DIC, institutions of similar size should treat that concentration of responsibility as an operational risk in its own right - the loss or unavailability of that one person represents a maintenance and, by extension, security gap with no fallback. Third, given that respondents linked system reliability directly to internet connectivity, institutions should evaluate prospective systems specifically for their tolerance of intermittent connectivity (offline modes, opportunistic sync) rather than assuming a stable, always-on connection that a resource-constrained setting may not reliably provide. Fourth, and specifically for vendors: a system marketed to institutions in low-connectivity regions should be engineered and tested for intermittent-connectivity operation as a core requirement, not an edge case, given that every connectivity-related complaint in this dataset originated from ordinary day-to-day use rather than an unusual outage.

Fifth, institutions should resist the temptation to treat a successful initial rollout as evidence that infrastructure planning is complete. DIC's data show respondents describing both early implementation difficulties and ongoing, later-stage performance problems, indicating that infrastructure adequacy should be reassessed on a recurring schedule - annually at minimum - rather than assumed to remain constant once a system has been successfully adopted. This recommendation follows directly from this paper's central argument: infrastructure is a continuous operating

condition, not a one-time hurdle, and institutional planning processes should be restructured to reflect that reality rather than continuing to treat infrastructure as a line item that is resolved once and then forgotten.

4.10 Toward a Minimum Viable Infrastructure Checklist

Table 3: Proposes a minimum viable infrastructure checklist for institutions operating under resource constraints similar to DIC's, following the same design logic as the checklists proposed in the companion papers [1,2]: each item addresses a specific gap observed directly in the interview data, rather than replicating an enterprise infrastructure-management framework a small institution is unlikely to sustain.

Gap observed	Minimum action
No device-compatibility testing	Test on actual staff/parent devices before rollout; budget a contingency fund
No backup schedule	Weekly backup at minimum, scheduled for a known-good-connectivity window
No maintenance ownership	Assign update/patch responsibility explicitly, with a named backup person
Single point of IT failure	Cross-train at least one additional staff member on basic system administration
No connectivity contingency	Identify and document an offline or degraded-mode fallback procedure
No cost-tracking beyond procurement	Track total cost of ownership annually, including ad hoc compatibility purchases

Minimum Viable Infrastructure Checklist Derived from DIC Findings

Each item is deliberately achievable without major capital investment, and the fourth item in particular - cross-training a backup staff member - directly addresses the single-point-of-failure risk identified in Section 3.3, converting a structural vulnerability into a manageable, low-cost mitigation.

It is worth being explicit about the checklist's sequencing logic, following the same approach used in the companion papers' checklists. The first two items - device-compatibility testing and a backup schedule - address root causes that, left unaddressed, will continue generating the reactive cost expansion and data-loss exposure documented in this paper regardless of how well the remaining items are executed. The third and fourth items address the single-point-of-failure risk discussed in Section 3.3, which this paper's findings suggest is the most structurally significant gap identified across all three papers in this research program, since it compounds governance, literacy, and infrastructure risk simultaneously in a single individual. The fifth item, a documented offline fallback procedure, is a direct response to the "system jams…needs strong internet" pattern reported by multiple respondents, converting an ad hoc, in-the-moment improvisation into a known, rehearsed procedure. The sixth item, annual total-cost-of-ownership tracking, is the least urgent in the short term but the most important over a system's multi-year operational life, since it is precisely the absence of this kind of tracking that allowed the mini-iPad purchase to occur as an unplanned surprise rather than an anticipated, budgeted contingency.

4.11 Designing an Infrastructure Resilience Metric

Paralleling the workaround-incidence metric proposed in the companion literacy paper [2], this section sketches a minimal infrastructure resilience metric appropriate for an institution with DIC's resource profile, since the absence of an existing lightweight standard for this purpose is itself a research gap this paper's argument implies needs filling.

A minimal resilience metric could combine three low-burden observations, collected monthly. First, a simple count of reported system-performance incidents (jams, failed scans, connectivity drop-outs) gathered by asking frontline staff to log any such incident during a defined observation period, rather than through automated instrumentation the institution does not currently have. Second, a binary check of whether the most recent scheduled backup completed successfully, which requires no new technology beyond a habit of checking and recording the outcome of an existing backup routine. Third, a simple tally of any ad hoc, unbudgeted technology purchases made in response to a compatibility or performance problem during the period, directly operationalizing the total-cost-of-ownership concern raised in Section 3.1 and providing the institution with a running record of exactly the reactive cost pattern this paper identifies as underappreciated in prior digitization-cost research.

As with the workaround-incidence metric proposed for the literacy dimension of this research program, none of these three measures requires specialized technical infrastructure, which is deliberate: a metric that itself requires the same technical sophistication and connectivity the institution is still working to establish would not be usable during the exact period it is meant to monitor. Tracked over successive months, a composite "infrastructure resilience index" combining

all three measures could help an institution distinguish a genuinely improving trajectory from a stable but fragile one, and could, with further validation across multiple institutions, provide researchers a standardized instrument for comparing infrastructure resilience across resource-constrained settings - addressing the specific comparability gap that Section 3.7 identifies as a limitation of this single-institution study.

Limitations

This paper's evidence on infrastructure and cost is drawn from stakeholder perceptions and anecdotes rather than measured system uptime, bandwidth logs, or itemized maintenance budgets, so the frequency and severity of the "system jam" pattern described by respondents cannot be quantified precisely from this dataset. The proposed AI-assisted directions in Section 3.4 are conceptual extensions motivated by the case-study findings rather than technologies piloted at DIC itself, and their feasibility for an institution with DIC's specific connectivity and staffing profile has not been tested. As with the companion papers, two citations reproduced from the original case-study report (Fan et al., 2018; Evens & Hauttekeete, 2011) lack complete bibliographic entries in that report's own reference list and are flagged accordingly in the references below. Finally, the comparative reading against enterprise CSPM and CWPP literature in Sections II and V is offered as a conceptual bridge rather than an empirical test; no CSPM or CWPP tooling was deployed or evaluated at DIC, and the applicability of enterprise-derived cost and effectiveness figures to a resource-constrained school setting remains an open empirical question.

A further limitation concerns generalizability beyond both the education sector and the Ghanaian setting examined here. DIC's specific combination of urban connectivity (better than many rural Ghanaian institutions would experience), a private-school budget model, and a diverse, parent-owned device ecosystem is a particular configuration of constraints that may not transfer directly to institutions in other sectors, other countries, or more rural settings. The core mechanism this paper identifies - that infrastructure and cost constraints recur throughout a system's operational life rather than concentrating at adoption, and that this recurrence has direct security consequences - is plausibly generalizable in structure, but the specific magnitude and character of the constraints (which devices are incompatible, how severe the connectivity gaps are, how large the reactive cost expansions turn out to be) would need to be re-examined for any other setting before this paper's specific findings could be assumed to transfer.

A final limitation concerns the infrastructure resilience metric proposed in Section 3.5: like the workaround-incidence metric proposed in the companion literacy paper, it is a conceptual proposal derived from this paper's findings rather than an instrument that has been piloted, validated, or tested for reliability. Before any institution or researcher adopts it, it would require the kind of validation work - testing against independently verified incident logs, piloting across more than one institutional setting - that this single-institution qualitative study was not designed to provide.

5 CONCLUSION

Divine International College's experience with PROCARE shows that infrastructure and cost constraints do not end at adoption; they recur as connectivity problems, device incompatibility, and an ongoing maintenance burden that a single, likely under-resourced IT function must absorb - with direct consequences for data security, not just system convenience. Institutions in similar settings should budget for infrastructure as a continuous operating cost, including a realistic device-compatibility contingency fund, rather than a one-time procurement line item. AI-assisted, low-bandwidth monitoring and maintenance tools, adapted downward from the enterprise CSPM, CWPP, and AI-incident-response literatures reviewed in Section 3.1, may offer a path to sustaining security-relevant upkeep without proportionally scaling IT staff, but - consistent with the companion papers in this research program [1,2] - this remains a direction for future empirical testing, not a validated solution ready for institutional adoption.

The broader implication of this paper, read alongside its two companions, is that resource-constrained institutions like DIC are not failing to meet a governance, literacy, or infrastructure standard designed with them in mind; they are being measured, implicitly, against standards and tools designed for better-resourced settings. Closing the gaps this three-paper research program identifies - in authentication and privacy governance, in onboarding and literacy support, and in infrastructure and maintenance capacity - will likely require the security, digital-preservation, and AI research communities to design deliberately for the resource-constrained case from the outset, rather than treating it as a simplified afterthought of enterprise-grade solutions.

For institutional leaders, the single most actionable recommendation from this study is to treat infrastructure and maintenance capacity as an explicit, recurring line in institutional risk assessment, rather than a background assumption revisited only when something breaks. The mini-iPad purchase documented in this case study is a useful diagnostic question any institution can ask of itself: has our technology budget ever had to absorb an unplanned compatibility or performance expense, and if so, was that expense tracked as part of the true cost of the system, or written off as an unrelated one-time cost? An institution that cannot answer this question is unlikely to have an accurate picture of what its digital preservation system actually costs to sustain.

For researchers, this paper's main contribution is the reframing of infrastructure and cost as an ongoing security condition rather than a one-time adoption barrier, and the proposal of a lightweight resilience metric, in Section 3.5, as a starting point for measuring that condition empirically rather than relying on the retrospective, anecdotal accounts this qualitative study necessarily depends on. For vendors serving resource-constrained institutions in low-connectivity regions, the implication is that intermittent-connectivity tolerance, device-compatibility breadth, and low-maintenance-burden design are not peripheral engineering nice-to-haves but core requirements directly tied to the security and continuity of the institutions such systems are meant to serve - a vendor that under-invests in these areas is, in effect, exporting infrastructure risk to every resource-constrained client institution it serves, in the same way the companion literacy paper argues under-invested onboarding design exports literacy risk downstream to client institutions least equipped to absorb it.

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] Y. Fan *et al.*, "Impact of internet connectivity on digital content transfer and preservation efficacy," as cited in the source case-study report, 2018. *[Full venue details not independently verified; reader should confirm before further citation.]*

[2] G. G. Dagher, J. Mohler, M. Milojkovic, and P. B. Marella, "Ancile: Privacy-preserving framework for access control and interoperability of electronic health records using blockchain technology," *Sustainable Cities and Society*, vol. 39, pp. 283–297, 2018.

[3] W. Y. Arms, *Digital Libraries*. Cambridge, MA: MIT Press, 2001.

[4] E. Lapworth, S. Jones, and M. Georgieva, "Microfilm, manuscripts, and photographs: A case study comparing three large-scale digitization projects," *Journal of Contemporary Archival Studies*, vol. 6, no. 1, 2019.

[5] O. Mosweu and D. Rakemane, "The role of records management in ensuring good governance in Africa," *Journal of the South African Society of Archivists*, vol. 53, pp. 103–123, 2020.

[6] T. Evens and L. Hauttekeete, "Challenges of digital preservation for cultural heritage institutions," as cited in the source case-study report, 2011. *[Full venue details not independently verified.]*

[7] M. A. Awamleh and F. Hamad, "Digital preservation of information sources at academic libraries in Jordan: An employee's perspective," *Library Management*, vol. 43, no. 1–2, pp. 172–191, 2022.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.