Cybersecurity Risks in DMZ-Hosted Examination Systems and Their Implications for Assessment Integrity: A Passive Cyber-Exposure Analysis

Christine Simfukwe*, Dennis Mulendema, Alice P. Shemi, Maybin Lengwe, Nchimunya Chaamwe

Copperbelt University

***Corresponding Author**

DOI: https://dx.doi.org/10.47772/IJRISS.2026.100800386

Received: 25 July 2026; Accepted: 30 July 2026; Published: 06 September 2026

ABSTRACT

The digital transformation of national examination systems has significantly improved service delivery across the education assessment sector. However, the increasing reliance on publicly accessible digital infrastructure has expanded the cyberattack surface, exposing examination systems to threats that jeopardize the integrity, confidentiality, and availability of assessment processes. The aim of the study was to assess the external cybersecurity exposure of national examination systems across Africa by identifying publicly observable vulnerabilities affecting examination infrastructure.

A systematic, non-intrusive cybersecurity exposure analysis was conducted on 15 national examination bodies across 10 countries in four African sub-regions. Using a passive reconnaissance methodology based on open-source intelligence (OSINT) tools, including Shodan, Censys, and Nmap, the study examined externally observable indicators of cybersecurity risk without intrusive testing, credential use, or system exploitation. Approximately 35 subdomains and 25 unique IP addresses were assessed, identifying 43 unique Common Vulnerabilities and Exposures (CVEs) with a mean vulnerability age of 8.4 years.

The findings reveal widespread and severe cybersecurity weaknesses. Approximately 80% of identified CVEs were rated Critical or High severity, 93% of institutions exhibited at least one exploitable publicly accessible vulnerability, and 47% were exposed to remote code execution vectors. Key risks included publicly exposed MySQL database ports, unpatched legacy protocols, outdated SSL/TLS configurations, and vulnerabilities dating from 1999 to 2025 that remained active in production environments. These weaknesses present credible threats to examination integrity, including unauthorized database access, candidate data exfiltration, and system compromise.

The study concludes that cybersecurity risks within African national examination systems are fundamentally governance challenges rather than isolated technical deficiencies. Strengthening cyber resilience requires the adoption of Zero Trust Architecture, structured patch management, rigorous network segmentation, and continuous vulnerability monitoring to safeguard the credibility and operational continuity of national examination systems.

Keywords: Cybersecurity; Examination Systems; Cybersecurity Governance

INTRODUCTION

The accelerated digitalisation of national examination systems (driven in large part by the operational pressures of the COVID-19 pandemic) has fundamentally altered the infrastructure profile of public sector assessment bodies worldwide. Services that were historically conducted entirely within controlled, paper-based environments, including candidate registration, examination administration, results dissemination, and certificate verification, are now delivered through publicly accessible digital platforms hosted on internet-facing

infrastructure [1][2]. This transition has yielded demonstrable improvements in efficiency, geographic reach, and administrative scalability. It has simultaneously introduced a category of systemic risk that conventional assessment integrity frameworks were not designed to address: infrastructure-level cybersecurity exposure.

Examination authorities operating digital platforms typically employ Demilitarised Zone (DMZ) network architectures to manage the tension between public accessibility and internal system protection. Publicly accessible examination services should be configured within a DMZ and appropriately isolated from internal networks hosting sensitive examination data, consistent with NIST guidance on secure server configuration [3]. When correctly configured, DMZ architectures provide meaningful network segmentation and limit the blast radius of external attacks. When misconfigured, inadequately maintained, or combined with unpatched software, they simultaneously maximise external accessibility and maximise external exposure creating a threat surface that is both broad and directly relevant to assessment integrity outcomes.

The scale of this threat is not theoretical. Lallie et al., analysing 58 cyberattacks affecting educational institutions, identified ransomware as the most prevalent external threat vector, with examination and student data representing primary targets [4]. Globally, the education sector consistently ranks among the most frequently attacked industries [27]. According to Bitsight Research the education sector accounted for 54.3% of all Known Exploited Vulnerability incidents across sectors [30], further the education sector was the largest impacted sector in the 2024 MOVEit breach (>50% of breached organisations) [31]. This pattern is attributable to the sensitivity of the data held, the relative underfunding of cybersecurity compared to other critical sectors, and the operational constraints on system downtime during examination periods [15]. Despite this context, the existing research literature is heavily concentrated on examination integrity concerns at the candidate level such as online cheating detection, AI-based proctoring, and LMS security rather than on the security of the examination infrastructure itself [5][6][7].

This study addresses that gap directly. It investigates the external cybersecurity exposure of DMZ-hosted examination systems across 15 institutions in 10 African countries using passive reconnaissance methodology, with three specific objectives: to identify and quantify externally observable vulnerabilities across examined institutions; to assess the potential impact of those vulnerabilities on assessment integrity; and to develop evidence-based recommendations for strengthening the security posture of examination infrastructure. Table 1 summarises the study scope.

Table 1: Study Scope Summary

Dimension	Scope
Institutions examined	15 examination bodies across 10 countries
Sub-regions covered	Southern Africa, East Africa, West Africa, Central Africa
Subdomains scanned	~35 publicly accessible subdomains
Unique IP addresses identified	25
CVEs identified	~43 unique CVEs across all institutions
Severity profile	~80% rated Critical or High
Vulnerability age range	1999–2025 (mean: ~8.4 years)
Methodology	Passive reconnaissance; non-intrusive; no system access

The study makes three principal contributions to the literature. First, it provides the first systematic, multi-institution quantified baseline of cybersecurity exposure across African examination infrastructure. Second, it demonstrates empirically that infrastructure-level vulnerabilities constitute direct and operationally feasible

threats to assessment integrity not peripheral concerns by mapping specific CVEs and port exposures to concrete result manipulation, data exfiltration, and system disruption scenarios. Third, it identifies systemic governance failure as the primary determinant of observed risk posture, a finding with significant implications for policy design and for the development of AI-driven adaptive cybersecurity frameworks for public sector networks.

LITERATURE REVIEW

Digital Transformation of Assessment Systems

The digitalisation of national examination systems in African contexts has accelerated substantially since 2020, driven by pandemic-induced operational necessity and longer-term efficiency imperatives [8]. E-assessment platforms enable remote evaluation at scale, improving administrative efficiency and geographic accessibility across primary and secondary education systems. However, implementation in developing country contexts remains constrained by infrastructural limitations including unreliable internet connectivity, heterogeneous digital infrastructure, and variable institutional digital literacy that shape both the deployment architecture and the security posture of examination platforms [9].

Digital results management systems, which automate the processing, storage, and dissemination of examination outcomes, represent a particular risk concentration point. These systems aggregate sensitive data candidate identifiers, raw marks, moderated grades, and certification records into centralised repositories that are simultaneously required to be accessible to authorised external users and protected from unauthorised access or modification [10]. Educational information systems of this type have become high-value targets for cyberattacks, necessitating the implementation of layered security controls including encryption at rest and in transit, role-based access control, and multi-factor authentication [11][12]. The extent to which these controls are implemented in practice within African examination bodies infrastructure is a central empirical question that this study addresses.

National examination systems occupy a unique position in the digital public sector landscape: they process data of exceptional societal sensitivity determining student progression, university admission, and professional qualification while operating under significant constraints on system downtime, configuration change, and security testing [13][14]. This combination of high-value data, operational constraint, and public accessibility creates a risk profile that is more acute than that of many other public sector digital systems and thus demands sector-specific empirical analysis.

Cybersecurity Threats in Examination Systems

The digitalisation of examination infrastructure has expanded the attack surface of national assessment systems across multiple threat categories. These include unauthorised access, data tampering, ransomware attacks, and distributed denial-of-service (DDoS) campaigns capable of disrupting examination processes and compromising sensitive candidate data at scale [15]. High-stakes examination systems are particularly attractive targets due to the critical nature of the data they process and the severe reputational, legal, and social consequences of successful compromise.

Examination result manipulation represents the most direct threat to assessment integrity and can be achieved through both external attack and insider threat vectors [16]. Weak authentication mechanisms, inadequate access controls, and insufficient monitoring of database activity increase the probability of such incidents. Critically, the exploitation of known vulnerabilities including outdated software components, misconfigured servers, and exposed service ports does not require sophisticated attack capabilities; many of the CVEs identified in this study have publicly available proof-of-concept exploit code and have been weaponised in documented attacks against comparable infrastructure [17].

Real-world incidents provide empirical grounding for these risk assessments. Ransomware attacks have disrupted institutional examination systems and exfiltrated sensitive data across multiple documented incidents in the education sector globally [4]. Examination system breaches resulting in unauthorised access to

examination materials and manipulation of candidate records have been reported in multiple African contexts, though systematic, institution-level exposure data for the region remains limited [13]. This study generates that empirical baseline.

DMZ Architecture in Public Examination Systems

A Demilitarised Zone (DMZ) is a network architectural pattern in which public-facing services are positioned in a logically isolated segment between the public internet and internal network tiers [18]. In national examination system deployments, DMZ environments typically host candidate-facing services registration portals, results verification systems, and certificate download platforms while internal networks retain sensitive data including raw examination records, grading databases, and administrative systems, as illustrated in Figure 1.

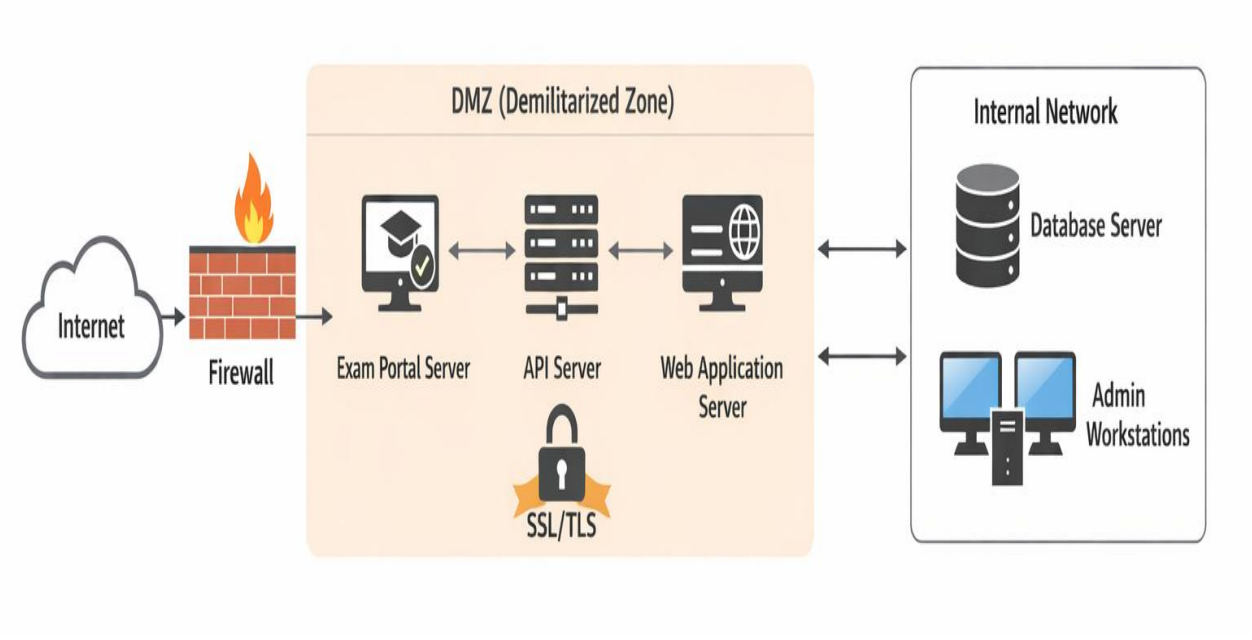

Fig. 1. Secure DMZ Architecture for Digital Examination Systems

The security value of a DMZ is conditional on correct configuration and ongoing maintenance. When properly implemented, DMZ architectures enforce network segmentation that limits the impact of external compromise: an attacker who gains control of a public-facing DMZ asset cannot, in principle, directly reach internal-tier systems without traversing additional security controls. However, the findings of this study demonstrate that in practice, a substantial proportion of examined institutions fail to enforce this segmentation effectively. The direct public accessibility of MySQL database ports (port 3306) on 40% of examined institutions a configuration that should be architecturally impossible in a correctly implemented DMZ illustrates the gap between the security model and its operational realisation. Misconfiguration, inadequate firewall rule management, and the deployment of legacy software within DMZ environments collectively undermine the protective architecture and convert it from a security control into a structured exposure surface.

Cybersecurity Frameworks for Protecting Assessment Systems

Established cybersecurity frameworks provide structured methodological approaches for managing institutional risk and protecting critical information systems. The NIST Cybersecurity Framework (CSF) organises security activities across five core functions Identify, Protect, Detect, Respond, and Recover which collectively constitute a lifecycle approach to cybersecurity risk management applicable to examination system environments [19]. Of particular relevance to the findings of this study are the Protect and Detect functions: the prevalence of unpatched CVEs with mean ages exceeding eight years, and the direct public exposure of database service ports, indicate that basic Protect-function controls patch management, network segmentation, and access control enforcement are not systematically implemented across the examined sector.

The CIA Triad (Confidentiality, Integrity, and Availability) provides the foundational risk classification framework for information security and maps directly onto the assessment integrity concerns of this study [19]. In examination system contexts, confidentiality failures enable unauthorised access to examination questions and candidate records; integrity failures enable result manipulation and data tampering; and availability failures disrupt examination delivery and results dissemination. The vulnerabilities identified in this study including remote code execution vectors (threatening all three dimensions), authentication bypass (threatening confidentiality and integrity), and cleartext data transmission (threatening confidentiality) represent systematic failures across all three CIA dimensions simultaneously.

Zero Trust Architecture (ZTA) extends beyond perimeter-based security models by enforcing continuous verification of all users and devices regardless of network location, assuming breach as a default posture, and applying least-privilege access controls at every system interaction [20] [28]. Applied to examination system environments, ZTA principles would eliminate several of the most critical risk pathways identified in this study: continuous verification requirements would neutralise authentication bypass CVEs; least-privilege enforcement would prevent direct database access from public-facing tiers; and breach assumption posture would mandate the monitoring and segmentation controls that are currently absent across the majority of examined institutions.[29]

Literature Gap Analysis

Table 3 maps the existing research literature across relevant domains against the gaps addressed by this study. The analysis demonstrates that while substantial research exists on candidate-level examination integrity concerns and general education sector cybersecurity threats, a systematic empirical analysis of cybersecurity exposure within national examination infrastructure (particularly in African contexts) is absent from the literature. This study addresses three primary gaps: it provides the first quantified multi-institution baseline for African examination bodies cybersecurity posture; it introduces CVE-level vulnerability profiling including age, severity, and institutional distribution analysis as a methodology for examination infrastructure risk assessment; and it empirically establishes the link between infrastructure-level vulnerabilities and specific assessment integrity compromise scenarios.

Table 3: Literature Coverage and Gap Analysis

Research Area	Existing Coverage	Gap Addressed by This Study
E-assessment quality and validity	Substantial — multiple systematic reviews	Not addressed; out of scope
Online examination integrity (cheating)	Substantial — AI proctoring, facial recognition, behavioural analytics	Not addressed; out of scope
LMS security (Moodle, Canvas, Blackboard)	Moderate — authentication, session management	Not addressed; out of scope
Ransomware in education sector	Moderate — Lallie et al. (2021) and sector reports	Partially addressed; this study extends to exam-specific infrastructure
Cybersecurity of national examination infrastructure	Minimal — no systematic empirical study identified	PRIMARY CONTRIBUTION: passive exposure analysis of DMZ-hosted national exam systems
African examination body cybersecurity posture	Absent from literature	PRIMARY CONTRIBUTION: first quantified multi-institution baseline for the African context

Research Area	Existing Coverage	Gap Addressed by This Study
CVE-level vulnerability profiling of exam systems	Absent from literature	PRIMARY CONTRIBUTION: CVE age, severity, and institutional distribution analysis

METHODOLOGY

This study employs a descriptive and exploratory research design using passive and minimally interactive reconnaissance techniques to assess the externally observable cybersecurity exposure of public-facing examination system infrastructure. The methodology was designed to generate empirically grounded, reproducible findings while adhering strictly to ethical research principles: no penetration testing, credential use, system exploitation, or unauthorised access was conducted at any stage. All findings are derived exclusively from publicly available and externally observable data.

Research Design and Ethical Boundaries

The study adopts a non-intrusive research design consistent with established ethical standards for passive cybersecurity research [23]. The analytical focus is the reconnaissance phase of the Cyber Kill Chain model the phase in which an attacker gathers externally observable information prior to exploitation as this phase yields information that is both ethically accessible and directly relevant to institutional risk posture [24]. By restricting analysis to this phase, the study generates findings that are representative of what a moderately capable adversary could discover without authorisation, thereby providing a realistic and conservative estimate of each institution's external attack surface.

Responsible disclosure principles were observed throughout. Where critical findings were identified that could present immediate risk to specific institutions, the study protocol required that such observations be communicable through appropriate channels without the public exposure of institution-identifying details. Accordingly, all institutions are anonymised in this publication using regional prefix codes (SAF, EAF, WAF, CAF) rather than their actual names or country identifiers.

CONCEPTUAL FRAMEWORK

The study's analytical framework integrates three methodological components drawn from the Cyber Kill Chain reconnaissance model [24]:

- Passive Reconnaissance: Collection of publicly available intelligence including OSINT data, DNS records, certificate transparency logs, and correlation with public vulnerability databases (NVD, CVE, Exploit-DB). This component generates the CVE profile and technology stack data that constitute the primary analytical dataset.

- Non-intrusive Active Reconnaissance: Limited interaction with public systems through standard HTTP/HTTPS requests, DNS queries, and controlled port observations. All interactions were limited to requests that any standard web browser or network diagnostic tool would generate in normal operation, ensuring no disruption to operational systems.

- Data Aggregation and Correlation: Integration of collected data to map domains, subdomains, exposed services, technology stacks, and CVE associations. Correlation with NVD severity ratings and Exploit-DB availability data enabled risk classification of identified exposures.

Data Collection Instruments and Scope

Data collection employed four primary instruments: Shodan [25], and Censys (internet-wide scanning indices providing port, service, banner, and certificate data for publicly observable IP addresses) [26]; DNS analytics tools (for subdomain enumeration and hosting geography identification); and purpose-built Python scripts (for automated CVE correlation against NVD records based on identified software versions and banners). The combination of these instruments enabled systematic coverage across all examined institutions without requiring direct system interaction beyond standard HTTP requests.

The study assessed approximately 35 publicly accessible subdomains across 15 institutions, identifying 25 unique IP addresses and cataloguing the technology stack, open port profile, TLS configuration, and CVE associations for each asset. The institution sample was constructed to provide geographic coverage across four African sub-regions Southern Africa (6 institutions), East Africa (3), West Africa (3), and Central Africa (3) while remaining constrained to institutions for which sufficient publicly observable infrastructure was identifiable.

Exposure Indicators and Classification Framework

The study assessed six categories of exposure indicator, selected on the basis of their direct relevance to assessment integrity risks and their observability through passive reconnaissance methods. Table 4 presents the indicator framework, the observable signal used to detect each indicator, its risk category, and its specific relevance to assessment integrity.

Table 4: Exposure Indicator Framework

Exposure Indicator	Observable Signal	Risk Category	Assessment Integrity Relevance
Firewall misconfiguration	Non-standard ports open; no filtering evidence	Network	Enables direct access to internal-tier services
Network segmentation weakness	Database ports (3306, 5432) publicly reachable	Architecture	Bypasses DMZ isolation; direct result DB access
Outdated TLS/SSL configuration	TLS 1.0/1.1 active; SSLv3 detected; no HTTPS	Encryption	Cleartext transmission of candidate PII and results
Weak authentication mechanisms	CVEs enabling auth bypass (e.g. CVE-2012-2122)	Access Control	Unauthorised database access without credentials
Exposed service ports	FTP (21), RDP (3389), Telnet (23), MySQL (3306)	Exposure	Direct attack surface for RCE, data theft, manipulation
Known CVEs in exposed systems	CVE matches via NVD/Exploit-DB correlation	Vulnerability	Confirmed exploitable risk in publicly observable systems

Each institution was assigned an overall risk classification (CRITICAL, HIGH, or MODERATE) on the basis of its composite exposure profile across these six indicator categories, weighted by the severity of identified

CVEs and the directness of the attack pathway to examination data assets. The classification methodology is consistent across all institutions, enabling a comparative benchmarking.

Data Analysis Approach

Collected data was structured into an asset inventory and exposure classification framework, with each scanned asset recorded against its institution, subdomain, IP address, open port profile, technology stack, TLS configuration, hosting geography, and associated CVE list. CVE records were enriched with NVD severity ratings, publication dates, and Exploit-DB availability indicators, enabling the vulnerability age and severity distribution analysis.

Analysis was conducted across three levels: asset-level (individual subdomain profiles); institution-level (composite risk classification and attack surface mapping); and sector-level (aggregate CVE distribution, technology stack prevalence, and cross-institutional vulnerability pattern analysis). Results are presented using descriptive statistics including counts, percentages, means, and distributional characterisations and structured tables to enable pattern identification and comparative assessment across institutions and sub-regions.

Ethical Considerations

The study was conducted in full compliance with established ethical principles for passive cybersecurity research. No credential use, system access, exploitation, or intrusive testing was conducted at any stage. All data collected was derived from publicly available sources and open-access intelligence databases. The analysis was performed within a controlled Kali Linux cybersecurity research environment configured to prevent any unintended active interaction with target systems.

Institution anonymisation was applied throughout data collection, analysis, and reporting. No institution name, country name, or personally identifiable information appears in this publication. The study design and ethical protocols were reviewed and approved in accordance with the authors' institutional research ethics requirements prior to data collection.

FINDINGS

The passive cyber-exposure analysis encompassed 15 examination bodies distributed across 10 countries and four African sub-regions (Table 5), assessing approximately 35 publicly accessible subdomains and identifying 25 unique IP addresses. Of the 15 institutions examined, 14 yielded detectable vulnerabilities across externally observable indicators; one institution (SAF-05) returned zero CVE detections, representing the single cleanest security profile in the dataset. The analysis focused exclusively on externally observable exposure within DMZ-hosted systems and provides a representative baseline of cybersecurity risk across African examination infrastructure.

Table 5: Profile of Anonymised Examination Bodies Across African Sub-Regions

SN	Institution Code	Country Code	Region	Subdomains Scanned
1	SAF-01	SAF-C1	Southern Africa	4
2	SAF-02	SAF-C2	Southern Africa	3
3	SAF-03	SAF-C3	Southern Africa	2
4	SAF-04	SAF-C4	Southern Africa	2
5	SAF-05	SAF-C5	Southern Africa	1
6	SAF-06	SAF-C6	Southern Africa	3
7	EAF-01	EAF-C1	East Africa	2

SN	Institution Code	Country Code	Region	Subdomains Scanned
8	EAF-02	EAF-C2	East Africa	3
9	EAF-03	EAF-C3	East Africa	4
10	WAF-01	WAF-C1	West Africa	5
11	WAF-02	WAF-C2	West Africa	4
12	CAF-01	CAF-C1	Central Africa	2
13	CAF-02	CAF-C2	Central Africa	3
14	WAF-03	WAF-C3	West Africa	1
15	CAF-03	CAF-C3	Central Africa	1

Vulnerability Landscape — Aggregate Summary

Overall Vulnerability Distribution by Severity

A total of approximately 40 unique CVEs were identified across all examined institutions, of which an estimated 80% were rated Critical or High severity (Table 2). Critical vulnerabilities alone accounted for approximately 45% of the total (n≈18), with High-severity CVEs contributing a further 35% (n≈14). Medium and unclassified CVEs collectively constituted the remaining 20% (n≈8). This distribution indicates a severely elevated aggregate risk posture, in which the substantial majority of identified vulnerabilities represent exploitation-ready exposures rather than theoretical or low-probability risks.

Table 6: Distribution of Identified Vulnerabilities by Severity Level

Severity	Count of Unique CVEs	Percentage	Risk Interpretation
Critical	~18	~45%	Immediate exploitation risk; highest priority remediation
High	~14	~35%	Significant exposure; remediation within 30 days
Medium	~3	~8%	Moderate risk; remediation within 90 days
Not Classified / Other	~5	~12%	Awaiting classification; treat as High pending review
Total	~40	100%	80% of CVEs rated Critical or High

Most Frequently Occurring CVEs Across All Institutions

Table 7 presents the eight most frequently observed CVEs across the examined institutions, ranked by institutional prevalence. CVE-2024-38476 and CVE-2025-23048 both Critical-rated Apache HTTP Server vulnerabilities were identified across 10 or more institutions, making unpatched Apache the dominant systemic risk within the African examination ecosystem. Taken together, the five Apache-related CVEs in Table 3 affect between 6 and 10 institutions each, indicating that Apache patching failures are not isolated to individual institutions but constitute a sector-wide vulnerability pattern. By contrast, MySQL-related CVEs (CVE-2012-2122 and CVE-2016-6662) affect four institutions each, representing a concentrated but geographically distributed risk cluster.

Table 7: Most Frequently Observed CVEs Across Examined Institutions

CVE ID	Affected Service	Severity	Institutions Affected (n)	Description
CVE-2024-38476	Apache HTTP Server	Critical	10+	SSRF/mod_rewrite exploitation
CVE-2025-23048	Apache HTTP Server	Critical	10+	HTTP desynchronization
CVE-2024-38474	Apache HTTP Server	Critical	8	SSI command injection
CVE-2024-38475	Apache HTTP Server	High	10	Source code disclosure via mod_rewrite
CVE-2023-25690	Apache HTTP Server	Critical	9	HTTP request smuggling
CVE-2022-36760	Apache HTTP Server	High	6	mod_proxy exploitation
CVE-2012-2122	MySQL	Critical	4	Authentication bypass
CVE-2016-6662	MySQL	Critical	4	Remote code execution

Open Port Analysis

Port Frequency Distribution Across All Scanned Assets

Table 8 summarises the frequency and risk classification of network services detected across the 35 scanned assets. Standard web ports (443 and 80) were near-universal, present on 91% and 86% of assets respectively. More critically, three service types with CRITICAL risk ratings were identified beyond standard web traffic: MySQL (port 3306) was publicly exposed on 7 assets (20% of the total), representing a fundamental architectural failure; FTP (port 21) was exposed on 5 assets (14%), enabling unauthenticated file access in combination with identified CVEs; and RDP (port 3389) was exposed on one asset (SAF-02), representing the highest single-asset risk given the presence of CVE-2019-0708 (BlueKeep). The co-exposure of these CRITICAL service ports with corresponding CVEs on the same or related assets substantially increases the probability of successful exploitation.

Table 8: Network Service Exposure and Associated Security Risk Levels

Port	Service	Assets Exposing (of 35)	% of Assets	Risk Level
443	HTTPS	~32	91%	Standard
80	HTTP	~30	86%	Moderate (unencrypted)
3306	MySQL	7	20%	CRITICAL
22	SSH	5	14%	High
21	FTP	5	14%	CRITICAL
25	SMTP	4	11%	High
8080	HTTP-Alt/Proxy	4	11%	High
3389	RDP	1	3%	CRITICAL

Port	Service	Assets Exposing (of 35)	% of Assets	Risk Level
135	MS RPC	1	3%	High
23	Telnet	Implied	—	CRITICAL
2082–2096	cPanel Ports	1	3%	High
5985	WinRM	1	3%	High

Institution-by-Institution Risk Assessment

Comparative Risk Matrix

Table 9 presents a comparative risk classification for all 14 institutions with detectable exposure. Of these, seven (50%) were classified as CRITICAL risk, five (36%) as HIGH, and two (14%) as MODERATE. No institution received a LOW or SECURE classification. The two MODERATE-rated institutions (SAF-03 and SAF-05) exhibited limited CVE profiles and no exposure of CRITICAL service ports, though both retained SSL/TLS weaknesses. The predominance of CRITICAL and HIGH classifications across geographically dispersed institutions indicates that the observed risk posture is systemic rather than attributable to isolated institutional failures.

Table 9: Institution-Level Cybersecurity Risk Matrix

Institution	CVEs Found	Critical CVEs	MySQL Exposed	FTP Exposed	RDP Exposed	SSL/TLS Issues	Overall Risk
SAF-01	8+ per subdomain	4	Yes (3 assets)	No	No	1 asset no SSL	CRITICAL
SAF-02	14	6	Implied	Yes	Yes	Yes	CRITICAL
SAF-03	3	0 (main)	No	Yes	No	Yes	MODERATE
SAF-04	6–8	2	Yes	No	No	Yes	HIGH
SAF-05	0 detected	0	No	Yes	No	Yes	MODERATE
SAF-06	5–7	3	No	No	No	Yes	HIGH
EAF-01	18	6	No	No	No	Yes	CRITICAL
EAF-02	18 per subdomain	6	No	No	No	Yes	CRITICAL
EAF-03	18+	6–8	No	No	No	Mixed	HIGH
WAF-01	8–9	5	Yes	No	No	Mixed	CRITICAL
WAF-02	5–6	4	No	No	No	Mixed	HIGH
CAF-01	8–9	5	Yes	Yes	No	Yes	CRITICAL
CAF-02	8–10	4	No	Yes	No	Yes	CRITICAL
WAF-03	5	4	No	No	No	TLS 1.2 only	HIGH

Selected High-Risk Case Analysis

The following cases represent institutions exhibiting the highest compound risk profiles in the dataset, selected to illustrate distinct vulnerability archetypes.

SAF-01 — CRITICAL

SAF-01 presents a compounding vulnerability profile distributed across multiple subdomains. MySQL port 3306 is publicly accessible on at least three assets, and two Critical MySQL CVEs — CVE-2012-2122 (authentication bypass) and CVE-2016-6662 (remote code execution) remain unpatched. The combination of exposed database ports and authentication bypass vulnerabilities on shared infrastructure creates a direct pathway for unauthorised database access without interaction with the application layer, enabling potential result manipulation or data exfiltration. The use of shared IP infrastructure across subdomains further increases the risk of cascading compromise across examination services.

SAF-02 — CRITICAL

SAF-02 exhibits the broadest individual attack surface in the dataset, with 14 CVEs spanning multiple service categories including FTP, RDP, TLS, and HTTPS. Critical vulnerabilities include CVE-2019-0708 (BlueKeep — wormable RDP remote code execution), CVE-2014-0160 (Heartbleed — TLS memory disclosure), and CVE-2016-0800 (DROWN — retrospective TLS decryption). The public exposure of RDP alongside BlueKeep renders this asset susceptible to unauthenticated remote takeover. Notably, a separate SAF-02 subdomain exhibits a secure configuration with no CVEs and port 443 only, demonstrating that secure configuration is technically achievable within the same institution and that the observed vulnerabilities reflect governance failure rather than infrastructural impossibility.

EAF-01 and EAF-02 — CRITICAL

Both EAF-01 and EAF-02 exhibit identical vulnerability profiles, each with 18 Apache-related CVEs replicated consistently across multiple subdomains. The homogeneity of the vulnerability profile across independent subdomains strongly suggests shared, unpatched server infrastructure and an estimated patch lag of eight or more years for the oldest identified CVEs. This pattern is consistent with a scenario in which a single server image or configuration template has been deployed without subsequent patching across all associated assets.

CAF-01 — CRITICAL

CAF-01 presents an overlapping exposure of legacy and current-era vulnerabilities across multiple service types. FTP anonymous access (CVE-1999-0197, 26 years old) coexists with MySQL authentication bypass (CVE-2012-2122) and remote code execution (CVE-2016-6662), alongside SSH exposure and FTP RCE (CVE-2015-3306). The co-presence of a 26-year-old unpatched vulnerability alongside more recent exposures is statistically consistent with a systemic absence of vulnerability lifecycle management across all service categories.

CAF-02 — CRITICAL

CAF-02 demonstrates a comparable pattern to CAF-01, with FTP-related vulnerabilities and SMTP spoofing risks contributing to a multi-vector attack surface. The coexistence of legacy FTP CVEs and more recent Apache vulnerabilities, combined with offshore hosting of the results platform in Canada, creates compounding risks spanning both technical and data governance dimensions.

Technology Stack Analysis

Web Server Distribution

Apache HTTP Server accounts for 51% of identified assets (n≈18 of 35) and is the source of all five most frequently occurring CVEs in the dataset (Table 10). The disproportionate CVE burden associated with Apache

— relative to Nginx (23%) and IIS (14%), which exhibited markedly lower CVE concentrations indicates that the choice of Apache alone is not the proximate cause of risk, but rather the systematic failure to apply available security patches to Apache deployments. This distinction is significant: the vulnerabilities identified are not zero-day exposures but publicly disclosed, patchable CVEs, many available for remediation for multiple years prior to this assessment.

Table 10: Web Server Distribution and Vulnerability Burden

Web Server	Assets (of ~35)	% of Assets	Vulnerability Implication
Apache httpd	~18	51%	Primary CVE source; 5 of the top 5 CVEs are Apache-related
Nginx	~8	23%	Lower CVE burden; generally better patched in this dataset
Microsoft IIS	~5	14%	Moderate exposure; requires Windows patching cadence
Cloudflare/CloudFront	~3	9%	Edge-cached; limited direct exposure
Other (OpenResty)	~1	3%	Insufficient sample for trend analysis

Database Exposure

The direct public exposure of database ports constitutes a fundamental violation of DMZ design principles, in which database servers should reside on internal network segments inaccessible from the internet. As shown in Table 11, MySQL port 3306 was found publicly accessible on seven assets across six institutions (40% of the dataset), with two additional institutions exhibiting MySQL-related CVEs that imply database service activity. No examination body operating a DMZ architecture should expose database ports directly to the internet; the observed exposure pattern is indicative of either misconfiguration or the complete absence of network segmentation controls.

Table 11: MySQL Database Exposure Across Examined Institutions

Database	Publicly Exposed Instances	% of Institutions	Institutions
MySQL (port 3306 open)	7 assets	40% (6 of 15)	SAF-01 (3), SAF-04 (1), WAF-01 (1), CAF-01 (2)
MySQL (CVEs, port inferred)	2 additional	13% (2 of 15)	SAF-02, CAF-01

SSL/TLS Protocol Analysis

Table 12 presents the SSL/TLS configuration profile across the 35 scanned assets. Approximately 77% of assets operate with acceptable or good TLS configurations (TLS 1.2+1.3 or TLS 1.3 only), indicating that encryption deployment is not uniformly absent across the sector. However, 12% of assets comprising those with no SSL/TLS detected and those with no port 443 — represent a Critical exposure category in which examination data, including candidate personally identifiable information and result transmissions, may be transferred in cleartext. A further 11% operate on TLS 1.2 only, which is approaching deprecation under current NIST guidance and should be prioritised for upgrade.

Table 12: SSL/TLS Configuration Profile Across Scanned Assets

TLS Configuration	Assets (of 35)	Percentage	Risk Assessment
TLS 1.3 only	12	34%	Good — current best practice
TLS 1.2 and TLS 1.3	15	43%	Acceptable — upgrade path clear
TLS 1.2 only	4	11%	Needs upgrade — NIST deprecating TLS 1.2
No SSL/TLS detected	2	6%	Critical — cleartext data transmission likely
Not applicable (no port 443)	2	6%	Critical — no encrypted channel available

Hosting and Geographic Distribution

Table 13 maps the geographic hosting profile of examined institutions. While 47% of institutions maintain at least their primary infrastructure within national borders, a significant proportion host examination-critical system including results platforms and candidate data repositories outside their national jurisdiction. Institutions hosting in the United States, United Kingdom, France, and Canada are subject to the legal and regulatory frameworks of those jurisdictions, which may conflict with national data protection legislation and create risk of compelled data access under foreign law. The pattern of Francophone African institutions (CAF-01, CAF-02) hosting in France, and West African institutions hosting in the UK, suggests that hosting decisions may be driven by historical or linguistic ties rather than data governance considerations.

Table 13: Geographic Hosting Distribution and Data Sovereignty Implications

Hosting Location	Institutions	% of Institutions	Data Sovereignty Risk
In-country	SAF-01, SAF-02, EAF-02, EAF-03, SAF-04, SAF-05, WAF-02 (main)	47%	Low — data sovereignty maintained
United States	WAF-01 (main), EAF-01	13%	High — subject to US jurisdiction/CLOUD Act
United Kingdom	WAF-01 (results), WAF-02 (3 subdomains), WAF-03	20%	High — GDPR/UK data law applies
South Africa	WAF-01 (digital certs), CAF-03, SAF-06 (Azure)	20%	Moderate — regional hub, POPIA applies
France	CAF-01 (2 subdomains), CAF-02 (main)	13%	High — subject to French/EU law
Canada	CAF-02 (xgestedu.com)	7%	High — results platform offshore

Vulnerability Age Analysis

A temporal analysis of the 43 unique CVEs identified across examined institutions reveals a distribution spanning 26 years (1999–2025), with an estimated mean vulnerability age of approximately 8.4 years. The distribution, summarised in Table 10, is right-skewed: while the most recent cohort (2023–2025) constitutes the largest single group (34.9%, n=15), legacy CVEs predating 2010 remain present in active production environments, indicating that remediation efforts are selective and reactive rather than systematic.

The two oldest cohorts (1999–2005 and 2011–2016) collectively account for 23.3% of all identified CVEs (n=10), all of which carry Critical or High severity ratings. This disproportionate concentration of high-severity vulnerabilities among the oldest cohorts suggests that prolonged non-remediation correlates with elevated institutional risk, as legacy CVEs tend to be extensively documented, publicly exploited, and straightforward to weaponise.

Table 14: Distribution of Vulnerability Age Across Examined Systems

Year Range	Unique CVEs (n)	% of Total	Mean Age (Years)	Observation
1999–2005	3	7.0%	22.5	FTP/Telnet vulnerabilities; 20–26 years unpatched
2011–2016	7	16.3%	11.2	MySQL, SSH, Heartbleed, DROWN
2017–2019	8	18.6%	7.0	Apache, SSH, BlueKeep
2020–2022	10	23.3%	4.0	Mixed service vulnerabilities
2023–2025	15	34.9%	1.0	Latest Apache CVEs; largest cohort
Total / Mean	43	100%	~8.4	Right-skewed; bimodal legacy/recent distribution

Table 15 provides a representative subset of legacy CVEs and their institutional distribution. Of the seven CVEs listed, six (85.7%) are rated Critical, with a mean vulnerability age of 13.4 years at the time of assessment. The institutional distribution of these legacy CVEs is particularly revealing: SAF-02 appears across six of the seven entries (85.7%), and the same four institutions (SAF-01, SAF-02, WAF-01, CAF-01) are consistently implicated in both MySQL-related CVEs. This recurrence pattern is statistically inconsistent with isolated or incidental patch failures and is instead consistent with systemic absence of a vulnerability lifecycle management function within those institutions. The presence of CVE-1999-0197 a 26-year-old FTP vulnerability enabling anonymous file access in a live production environment alongside 2025-era Apache CVEs provides particularly strong evidence that remediation, where it occurs, is neither comprehensive nor sustained.

Table 15: Representative Legacy Vulnerabilities and Affected Institutions

CVE	Year	Age (Years)	Severity	Impact	Institutions Affected (n)	Institutions
CVE-1999-0197	1999	26	Critical	FTP anonymous access	3	SAF-02, CAF-01, CAF-02
CVE-2005-0469	2005	20	High	Telnet buffer overflow	1	SAF-02
CVE-2012-2122	2012	13	Critical	MySQL auth bypass	4	SAF-01, SAF-02, WAF-01, CAF-01
CVE-2014-0160	2014	11	Critical	Heartbleed — memory disclosure	1	SAF-02
CVE-2016-0800	2016	9	Critical	DROWN — TLS decryption	1	SAF-02
CVE-2016-6662	2016	9	Critical	MySQL RCE	4	SAF-01, SAF-02, WAF-01, CAF-01
CVE-2019-0708	2019	6	Critical	BlueKeep — wormable RDP RCE	1	SAF-02

CVE	Year	Age (Years)	Severity	Impact	Institutions Affected (n)	Institutions
Mean / Total	—	13.4	—	—	2.1 avg	SAF-02 appears in 6 of 7 (86%)

Attack Surface Categorisation

Table 16 categorises identified vulnerabilities by attack vector, quantifying the proportion of institutions exposed to each vector and its potential impact on assessment integrity. Remote code execution vectors encompassing Apache, MySQL, FTP, and RDP vulnerabilities affect the largest share of institutions (47%, n=7), representing the highest-order threat to examination system availability and integrity. Direct database access vectors affect 27% of institutions (n=4) and present the most direct pathway to result manipulation. The remaining vectors authentication bypass, cleartext interception, traffic decryption, and phishing collectively create a multi-dimensional attack surface that compounds the overall risk posture of the sector.

Table 16: Attack Surface Categorisation and Associated Risks to Assessment Integrity

Attack Vector	Institutions Exposed (n)	% of Institutions	Potential Impact on Assessment Integrity
Direct Database Access (MySQL 3306)	4	27%	Result manipulation, candidate data theft, grade alteration
Remote Code Execution (Apache, MySQL, FTP, RDP)	7	47%	Complete server compromise, ransomware, data destruction
Authentication Bypass (CVE-2012-2122)	4	27%	Unauthorised database access without credentials
Cleartext Interception (HTTP-only)	2	13%	Candidate PII exposure, session hijacking
Remote Desktop Takeover (RDP/BlueKeep)	1	7%	Complete Windows system takeover, lateral movement
File System Access (FTP anonymous)	3	20%	Examination paper theft, malicious file upload
Encrypted Traffic Decryption (Heartbleed, DROWN)	1	7%	Historical and real-time traffic decryption
Phishing/Spoofing (SMTP)	2	13%	Targeted phishing of candidates and staff

Comparative Benchmarking

Institutions with Best Security Posture

Table 17 identifies the four highest-performing assets in the dataset. All four share three common characteristics: zero detected CVEs, minimal port exposure limited to standard web ports, and current TLS configurations. Notably, three of the four are subdomains or secondary assets of institutions that also host CRITICAL-rated

assets (SAF-02 and SAF-06) demonstrating that secure configuration is operationally achievable within the same institutional environments that also exhibit critical vulnerabilities. This intra-institutional variation is itself a significant finding, suggesting that the determinant of security posture is governance and configuration management rather than resource availability.

Table 17: Top-Ranked Assets by Security Posture

Rank	Institution	Justification
1	SAF-03 (main site)	No CVEs detected; Nginx/OpenResty stack; TLS 1.2+1.3; minimal port exposure
2	SAF-02 (subdomain A)	No CVEs detected; IIS on Windows Server; port 443 only; TLS 1.2+1.3
3	SAF-02 (subdomain B)	No CVEs; Microsoft Azure-hosted; ports 80/443 only; TLS 1.2+1.3
4	SAF-06	No CVEs detected; minimal port exposure; TLS 1.3

Institutions with Worst Security Posture

Table 18 identifies the five lowest-performing institutions in the dataset. SAF-02 ranks lowest overall despite hosting some of the dataset's best-configured subdomains, reflecting the most extreme intra-institutional security variance observed. EAF-01 and EAF-02 are distinguished by the highest raw CVE counts (18 each), with identical vulnerability profiles suggesting shared unpatched infrastructure. CAF-01 is notable for the co-presence of the oldest CVE in the dataset (CVE-1999-0197, 1999) alongside contemporary vulnerabilities, representing the most extreme example of patch management failure by age distribution.

Table 18: Lowest-Ranked Institutions by Security Posture

Rank	Institution	CVEs (n)	Critical CVEs	Justification
1	SAF-02	14	6	BlueKeep, Heartbleed, DROWN; FTP and RDP directly exposed; broadest attack surface
2	SAF-01	8+ per subdomain	4	MySQL exposed on 3 assets; auth bypass + RCE unpatched; no SSL on service portal
3	EAF-02	18 per subdomain	6	Identical CVEs across 3 subdomains; estimated 8+ years without patch application
4	EAF-01	18	6	Identical unpatched Apache profile; shared infrastructure with EAF-02 suspected
5	CAF-01	8–9	5	FTP anonymous access; MySQL exposed; SSH; 8 critical/high CVEs including legacy

Key Findings, Implications, and Recommended Actions

Table 19 consolidates the principal findings of this analysis, their implications for assessment integrity, and prioritised remediation recommendations. Actions are classified by urgency: IMMEDIATE actions address exploitable vulnerabilities with direct pathways to assessment system compromise; URGENT actions address sector-wide systemic risks; SHORT-TERM and MEDIUM-TERM actions address architectural and governance

deficiencies; and ONGOING actions address structural security maturity gaps requiring sustained institutional commitment.

Table 19: Summary of Key Findings, Implications, and Recommended Actions

Priority	Finding	Description	Implications for Assessment Integrity	Recommended Action
IMMEDIATE	Database Ports Exposed	MySQL port 3306 exposed on 7 assets (20%) with authentication bypass and RCE CVEs	Direct unauthorised modification of examination results	Close database ports to public internet; enforce strict firewall rules and access controls
IMMEDIATE	Critical Protocol Exposure	FTP exposed on 5 assets (14%); RDP on 1 asset; Telnet implied	Enables unauthorised access, ransomware, and large-scale data breaches	Disable legacy protocols; enforce SSH, SFTP, and RDP gateways with MFA
IMMEDIATE	Absent or Weak Encryption	12% of assets lack SSL/TLS or have no HTTPS capability	Candidate PII and examination results may be transmitted in cleartext	Deploy TLS 1.3; enforce HTTPS-only via HSTS; deprecate TLS 1.2 where feasible
URGENT	Systemic Apache Vulnerabilities	88% of institutions exhibit unpatched Apache CVEs across multiple versions	Increased risk of RCE, SSRF, and request smuggling across the majority of institutions	Update Apache to current release; deploy WAF; implement automated patch monitoring
URGENT	Legacy Vulnerabilities	43 CVEs spanning 26 years (mean age: 8.4 years); 7.0% predating 2006	Persistent exposure to well-documented, trivially weaponisable attack vectors	Implement formal patch management programmes with mandatory SLAs by severity
SHORT-TERM	Infrastructure Reuse	Multiple subdomains sharing IP addresses observed across institutions	Single compromise can propagate across multiple examination services	Improve network segmentation; isolate critical examination systems
MEDIUM-TERM	Offshore Hosting	53% of institutions host at least one service outside national jurisdiction	Data sovereignty and regulatory compliance risks under national data protection frameworks	Develop and enforce data sovereignty policies; assess contract obligations
ONGOING	Inconsistent Security Maturity	Security posture varies within institutions; some subdomains well-secured while others are critical	Uneven protection creates exploitable gaps within the same institutional environment	Standardise cybersecurity policies; conduct continuous automated security assessment

DISCUSSION

The findings of this study establish a statistically significant and operationally consequential relationship between cybersecurity exposure and the integrity of digitalised examination systems across African institutions. Across 15 institutions examined, 93% (n=14) exhibited at least one Critical or High severity vulnerability in an externally observable, publicly accessible system; 80% of the approximately 40 unique CVEs identified were rated Critical or High; and seven institutions (47%) were exposed to at least one remote code execution vector capable of enabling complete server compromise. These figures collectively indicate that the risk to assessment integrity is not marginal or theoretical it is structural and widespread.

Table 20 maps the principal finding categories identified in this study to their corresponding assessment integrity risks and relative magnitudes. The most direct threats to examination integrity were direct database exposure

and authentication bypass which affected 27% to 40% of institutions. These may create pathways for unauthorised modification of results that do not require sophisticated attack capabilities. CVE-2012-2122, for instance, permits MySQL authentication bypass using a known exploit that has been publicly available since 2012; an attacker with basic technical knowledge could leverage this vulnerability to access examination databases on any of the four affected institutions without valid credentials. The persistence of such vulnerabilities, 13 years after public disclosure, is not attributable to technical complexity but to the absence of governance mechanisms capable of enforcing remediation.

Table 20: Mapping of Key Findings to Assessment Integrity Risks

Finding Category	% of Institutions Affected	Primary Assessment Integrity Risk	Risk Magnitude
Critical/High severity CVEs (80% of total CVEs)	93% (14 of 15)	System compromise enabling result manipulation or data destruction	Critical
Direct database exposure (MySQL port 3306)	40% (6 of 15)	Unauthorised read/write access to examination results without authentication	Critical
Remote code execution vectors (Apache, MySQL, FTP, RDP)	47% (7 of 15)	Complete server takeover; ransomware; examination service disruption	Critical
Absent or degraded SSL/TLS	23% of assets	Cleartext transmission of candidate PII, session tokens, and results	High
Legacy protocols (FTP, Telnet, RDP)	40% (6 of 15)	Unauthenticated file access; network traffic interception	High
Offshore hosting of examination data	53% (8 of 15)	Foreign legal jurisdiction over candidate data; compelled disclosure risk	High
Intra-institutional security variance	Observed in 3 of 15	Inconsistent protection creates exploitable gaps within same environment	Moderate

A critical observation emerging from the data is the distinction between technical capacity and governance will. The intra-institutional variance identified at SAF-02 where a subdomain with zero CVEs and port-443-only exposure coexists alongside an asset with 14 CVEs including BlueKeep and Heartbleed demonstrates that the infrastructure and technical knowledge required to achieve secure configuration are present within the same organisation. This pattern, also observable at SAF-06, suggests that the determinant of security posture is not resource availability but institutional governance: specifically, whether security policies are consistently enforced across all systems within an institution's control. Where policy enforcement is inconsistent, secure islands within an otherwise vulnerable environment provide limited systemic protection, as an attacker can simply pivot to the least-secured asset.

Apache HTTP Server vulnerabilities emerged as the most prevalent software weakness identified in this study. They were present in 88% of the institutions analysed, and the two most common CVEs each affected ten or more institutions. This pattern indicates that the vulnerabilities are systemic rather than isolated, pointing to common deficiencies in software lifecycle management and patch governance across examination institutions. Apache httpd is an open-source project with a well-maintained public vulnerability disclosure and patch release process Further, all CVEs identified in this study had available patches at the time of assessment. The sector-wide failure to apply these patches is therefore not attributable to vendor non-disclosure or patch unavailability, but to the absence of automated patch monitoring, testing, and deployment pipelines across the majority of

examined institutions. This represents a systemic governance failure rather than a set of isolated technical oversights.

Table 21 presents a structured analysis of the governance dimensions implicated by the findings. Across six governance dimensions patch management, network segmentation, configuration standards, intra-institutional policy enforcement, protocol lifecycle management, and data sovereignty the observed evidence consistently points to the absence of formalised, enforceable institutional processes rather than to technical barriers. The mean CVE age of 8.4 years, combined with a maximum age of 26 years, is particularly diagnostic: it implies that vulnerability management, where it exists, operates as an ad hoc or reactive function rather than a continuous, policy-driven process integrated into the operational lifecycle of examination systems.

Table 21: Governance Gap Analysis — Evidence and Implications

Governance Dimension	Observed Evidence	Implication
Patch management	Mean CVE age of 8.4 years; CVEs spanning 26 years in active systems	No formalised patch lifecycle or remediation SLA in place
Network segmentation	MySQL port 3306 publicly exposed on 20% of assets	DMZ design principles not enforced; database tier not isolated
Configuration standards	Identical 18-CVE profile across EAF-01 and EAF-02 subdomains	No secure baseline configuration or post-deployment hardening applied
Intra-institutional policy enforcement	SAF-02 hosts both most vulnerable and best-configured assets	Policies exist but are applied inconsistently across systems
Protocol lifecycle management	CVE-1999-0197 (FTP anonymous access) active in 2025	No protocol deprecation or technology refresh programme
Data sovereignty governance	53% of institutions host data outside national jurisdiction	No enforceable data localisation policy or hosting governance framework

These findings carry important implications for how assessment integrity should be conceptualised in the context of digital examination systems. Conventional frameworks for assessment integrity that are primarily concerned with candidate misconduct, examiner bias, and process fairness, do not adequately account for infrastructure-level vulnerabilities as a vector for large-scale integrity compromise. The exposure of MySQL databases to unauthenticated internet access, combined with the presence of remote code execution vulnerabilities across nearly half of examined institutions, creates conditions under which examination results could be altered, deleted, or fabricated at scale, without any candidate-level action. This represents a qualitatively different threat model from traditional assessment integrity concerns and requires a correspondingly different institutional response.

The study further contributes to a growing body of evidence that cybersecurity in public sector digital infrastructure in developing country contexts is characterised by prolonged technical debt, governance fragmentation, and the prioritisation of functional deployment over security assurance. The co-presence of 1999-era FTP vulnerabilities and 2025 Apache CVEs within the same institutional environments is consistent with a pattern in which systems are deployed, connected to the internet, and subsequently left to accumulate exposure without structured review. Addressing this pattern requires not only technical remediation such as patching, segmentation, and protocol deprecation but the establishment of institutional cybersecurity governance frameworks. These will treat vulnerability management as an operational function with defined accountability, resourcing, and performance metrics.

The adoption of Zero Trust Architecture principles, which assume that no network segment or user is inherently trustworthy and enforce continuous verification at every access point, offers a structurally appropriate response to the DMZ-level vulnerabilities identified in this study. Similarly, robust network segmentation enforcing strict separation between public-facing web tiers, application tiers, and database tiers would eliminate the most direct risk pathway identified i.e. the public accessibility of MySQL databases. These architectural approaches are well-documented, widely implemented in comparably resourced public sector contexts, and do not require proprietary or high-cost technologies. Their absence in the majority of examined institutions reflects a governance gap rather than a technical or resource constraint.

Limitations Of Study

This study is subject to five principal methodological limitations, each of which introduces a specific form of bias or constraint on the generalisability of the findings. Table 22 presents a structured assessment of each limitation, its nature, the direction of bias it introduces, and its proposed mitigation in future research.

Table 22: Methodological Limitations, Bias Direction, and Mitigation Strategies

Limitation	Nature	Direction of Bias	Mitigation in Future Work
Passive reconnaissance only	No authenticated or intrusive testing; internal network not assessed	Underestimation — true vulnerability count likely higher	Controlled penetration testing with institutional consent
Dependency on indexed data sources	Shodan and Censys may not reflect real-time configurations	Potential misclassification — stale data may not reflect recent patches	Combine with active DNS and certificate transparency monitoring
Exploitability not confirmed	CVEs identified as exposure indicators, not verified exploits	Conservative — findings represent potential rather than confirmed risk	Authorised exploit validation in controlled environments
Sample size (n=15)	15 institutions across 10 countries; not a random sample	Selection bias possible; findings may not generalise to all African examination bodies	Expand to stratified sample of 50+ institutions
Single time-point assessment	Snapshot methodology; dynamic environments may change post-assessment	Temporal bias — some vulnerabilities may have been patched after scanning	Longitudinal tracking with quarterly reassessment

The most significant limitation is the exclusive reliance on passive reconnaissance methods. Because the analysis was confined to publicly observable indicators, open ports, banner information, certificate data, and CVE correlation vulnerabilities within authenticated systems, internal networks, or application logic layers could not be assessed. This introduces a systematic underestimation bias: the true vulnerability counts for each institution are likely higher than those reported, as internal systems may carry additional unpatched CVEs not observable from the public internet. Consequently, institutions classified as MODERATE risk on the basis of low observable CVE counts (SAF-03, SAF-05) may harbour additional exposures not captured by this methodology.

A second limitation relates to the temporal accuracy of the data sources employed. Shodan and Censys index public-facing assets at varying intervals, and their records may not reflect configuration changes made in the days or weeks immediately preceding the assessment window. This introduces a potential misclassification risk

in both directions: a recently patched system may still appear vulnerable in indexed data, or a newly exposed system may not yet be indexed. While this limitation is inherent to all passive reconnaissance methodologies and cannot be fully eliminated, it can be partially mitigated through cross-referencing multiple data sources and scheduling assessment windows to coincide with known indexing cycles.

Third, this study cannot confirm the exploitability of identified CVEs within the specific configuration context of each institution. CVE presence indicates that a vulnerable software version is likely deployed, but does not confirm that the precise conditions required for exploitation (such as specific module configurations, network routing rules, or authentication bypass preconditions) are met. The findings therefore represent exposure indicators and potential risk pathways rather than confirmed exploitable vulnerabilities. This distinction is important for interpreting the findings proportionately. The study identifies where risk is concentrated and where remediation should be prioritised, rather than providing a confirmed exploit inventory.

Finally, the sample of 15 institutions across 10 countries, while providing broad geographic coverage, was not drawn from a randomised sampling frame and cannot be assumed to be statistically representative of the full population of African examination bodies. Selection was constrained by the availability of publicly observable digital infrastructure and the scope of the research mandate. Findings should therefore be interpreted as indicative of sector-wide patterns rather than as precise population estimates. The observed prevalence figures should be regarded as lower-bound estimates subject to revision as the sample is expanded in future research.

Future Work

This study establishes a quantified passive exposure baseline for African examination institution cybersecurity and, in doing so, surfaces a set of empirical gaps that define a clear and urgent future research agenda. Table 23 summarises six research directions derived directly from the limitations and findings of this study, mapped to their proposed methodologies and anticipated contributions.

Table 23: Future Research Agenda — Directions, Methods, and Contributions

Research Direction	Methodology	Expected Contribution
Authorised penetration testing	Permission-based active testing with national examination authorities	Confirm exploitability of identified CVEs; validate passive findings
Longitudinal vulnerability tracking	Quarterly passive rescanning of same institutions over 2–3 years	Measure remediation velocity and patch management improvement over time
Expanded geographic sample	Stratified sample of 50+ institutions across all African sub-regions	Test generalisability of findings; identify regional variation in security posture
AI-driven threat detection	Machine learning models trained on examination system exposure data	Enable real-time anomaly detection and predictive vulnerability identification
Governance framework development	Qualitative interviews with examination authorities and IT security leads	Develop context-specific cybersecurity governance frameworks for the sector
Data sovereignty impact assessment	Legal and technical analysis of offshore hosting arrangements	Quantify regulatory risk and develop national data localisation policy templates

The most immediate near-term priority is the validation of passive findings through authorised, permission-based penetration testing conducted in formal collaboration with national examination authorities. Such testing would confirm the exploitability of the CVEs identified in this study within the specific configuration context of each institution, provide institution-specific remediation evidence, and generate the controlled empirical data needed

to quantify the actual risk reduction achievable through targeted patching. This phase of research would also enable the development of a sector-specific vulnerability severity scoring framework that accounts for the particular sensitivity of examination data and the assessment integrity consequences of different exploitation scenarios.

A longitudinal tracking study, conducting quarterly passive rescans of the same 15 institutions over a 24 to 36-month period, would provide the first time-series evidence of remediation velocity within African examination infrastructure. Such data would enable quantitative assessment of whether current findings reflect a static state of neglect or a slowly improving trend, and would generate the empirical basis for evaluating the impact of governance or policy interventions implemented in response to this or related research. Expansion of the geographic sample to a stratified set of 50 or more institutions would simultaneously strengthen the generalisability of the findings and enable statistical testing of hypotheses regarding regional variation in security posture.

The most strategically significant direction for future research is the development and empirical validation of an AI-driven adaptive cybersecurity framework specifically designed for large-scale public sector networks. The findings of this study provide the empirical motivation and sector-specific vulnerability profile that the proposed framework must address. Specifically, the study identified prolonged patch debt (mean CVE age: 8.4 years), heterogeneous technology stacks, inconsistent governance practices, and highly sensitive information assets whose integrity is critical to examination processes. Collectively, these findings indicate that static, rule-based security controls are insufficient to detect and respond effectively to the dynamic, multi-vector cyber threats facing examination systems. This limitation motivates ongoing doctoral research by the authors into AI-Driven Adaptive Cybersecurity for Large-Scale Public Sector Networks, which develops a systematic literature review and application-oriented framework for deploying machine learning-based threat detection, autonomous anomaly response, and adaptive policy enforcement within public sector network environments. The passive exposure baseline established in this study comprises the CVE age distribution, port exposure profile, and institutional risk stratification. These findings provide the primary empirical foundation for the proposed framework. Consequently, the framework's design is informed by the documented operational realities of African examination infrastructure instead of relying solely on generalized threat models. Future publications will explicitly extend and build upon the findings reported here, providing both the theoretical architecture and the applied validation that the present study's scope does not permit.

CONCLUSION

This study demonstrates, through systematic passive cyber-exposure analysis of 15 African examination institutions, that the digitalisation of public examination systems has introduced substantial and largely unmitigated cybersecurity risks to assessment integrity. The aggregate findings are unambiguous in their severity: 93% of examined institutions exhibited detectable critical or high-severity vulnerabilities; 80% of the approximately 40 unique CVEs identified were rated Critical or High; 47% of institutions were exposed to at least one remote code execution vector; and the mean age of identified vulnerabilities was approximately 8.4 years, with the oldest dating to 1999. These figures collectively indicate that the cybersecurity posture of the examined sector does not meet the baseline standards required to protect systems that underpin national credentialing and public trust in educational outcomes.

The study makes three principal contributions to the literature. First, it establishes a quantified baseline of cybersecurity exposure across African examination infrastructure using a reproducible passive reconnaissance methodology. This provides a comparative reference point for future longitudinal and cross-regional research. Second, the study provides empirical evidence that infrastructure-level vulnerabilities represent a direct and operationally feasible threat to assessment integrity. Rather than being merely theoretical or peripheral concerns, these vulnerabilities have practical security implications. This is demonstrated by mapping specific CVEs and exposed network ports to plausible attack scenarios involving result manipulation, data exfiltration, and system disruption. Third, it identifies systemic governance failure, rather than technical incapacity or resource scarcity, as the primary explanatory factor for the observed risk posture, a distinction with significant implications for the design of policy interventions.

The findings indicate that effective remediation of the identified risks does not require novel technologies or extraordinary resources. The vulnerabilities identified are, without exception, publicly disclosed, patchable, and addressable through well-established security engineering practices: timely application of vendor patches, enforcement of network segmentation between public-facing and database tiers, deprecation of legacy protocols, and consistent application of TLS 1.3 across all public-facing assets. What is absent is not technical knowledge but institutional governance: formalised patch management programmes with defined remediation SLAs by severity; configuration standards enforced consistently across all institutional assets; and accountability structures that treat cybersecurity as a continuous operational function rather than a one-time implementation exercise.

Ultimately, the integrity of national examinations in the digital age is inseparable from the security of the underlying digital infrastructure. Weaknesses in examination infrastructure directly undermine the fairness, confidentiality, authenticity, and availability of examination processes and outcomes.

The findings of this study illustrate the seriousness of this challenge that represents significant risks to the credibility and trustworthiness of national examination systems. Addressing these risks requires a holistic and sustained institutional commitment. Cybersecurity governance must be integrated into the core operational mandate of examination authorities, while national data protection policies should be aligned with secure examination hosting practices. Above all, the security of examination infrastructure should be recognised as a matter of public interest and afforded the same strategic importance as the examination process itself.

REFERENCES

1. L. Moccozet, A. Camacho, A. Bowman, O. Benkacem and P. Roth, "A Study of the Evolution of E-Assessment Practices Following the COVID-19 Pandemic in Higher Education," 2024 21st International Conference on Information Technology Based Higher Education and Training (ITHET), Paris, France, 2024, pp. 1-8, doi: 10.1109/ITHET61869.2024.10837598.
2. B. Taruberekera, "ZIMSEC Online Results Checking System: A Parent-Guardian Perception Survey," Institute of Education Conference. [Online]. Available: https://ojs.zou.ac.zw/index.php/iec/article/view/157
3. Scarfone, K., Jansen, W., & Tracy, M. (2008). Guide to general server security (NIST Special Publication 800-123). National Institute of Standards and Technology. Available: https://csrc.nist.gov/publications/detail/sp/800-123/final
4. Lallie, H. S., Thompson, A., Titis, E., & Stephens, P. (2024). Understanding cyber threats against universities, colleges, and schools. Computers, 14(2), 49. https://doi.org/10.3390/computers14020049
5. Newton, P., & Essex, K. (2022). How common is cheating in online exams and did it increase during the COVID-19 pandemic? A Systematic Review. Research Square [preprint]. https://doi.org/10.21203/rs.3.rs-2187710/v1
6. Noorbehbahani, F., Mohammadi, A., & Aminazadeh, M. (2022). A systematic review of research on cheating in online exams from 2010 to 2021. Education and Information Technologies, 27, 8413–8460. https://doi.org/10.1007/s10639-022-10927-7
7. Barut Tuğtekin, E. (2023). Scrutinizing Learning Management Systems in Practice: An Applied Time Series Research in Higher Education. The International Review of Research in Open and Distributed Learning, 24(2), 53–71. https://doi.org/10.19173/irrodl.v24i2.6905
8. Namatende-Sakwa, L., Lewinger, K., & Langsford, D. (Eds.). (2023). COVID-19 and education in Africa: Challenges, possibilities and opportunities. Routledge.
9. Mahlangu, G., & Makwasha, L. (2023). Factors affecting the adoption and use of online assessment for learning at Polytechnics in Zimbabwe. Cogent Education, 10(1), 2177475. https://doi.org/10.1080/2331186X.2023.2177475
10. Cremer, F., Sheehan, B., Fortmann, M., Kia, A. N., Mullins, M., Murphy, F., & Materne, S. (2022). Cyber risk and cybersecurity: A systematic review of data availability and cyber risk management. The Geneva Papers on Risk and Insurance — Issues and Practice, 47, 698–736. https://doi.org/10.1057/s41288-022-00266-6

11. Garg, M., & Goel, A. (2022). A systematic literature review on online assessment security: Current challenges and integrity strategies. Computers & Security, 113, 102544. https://doi.org/10.1016/j.cose.2021.102544 Garg, M., et al. (2022). Online assessment security: A systematic review. Computers & Security.
12. Guillen-Gamez, F. D., Tomczyk, Ł., Ruiz-Palmero, J., & Connolly, C. (2024). Digital security in educational contexts: Digital competence and challenges for good practice. *Computers in the Schools, 41*(2). https://doi.org/10.1080/07380569.2024.2390319
13. Yusuf Lawal, Modupeola Adepoju, & Ademeso Tosin Success. (2025). E-Examinations and national security: Challenges and prospects. IOSR Journal of Business and Management, 27(8), 09–16. doi: 10.9790/487X-2708050916
14. Oguguo, B. C., & Ocheni, C. A. (2023). Cybersecurity: A tool for curbing examination breaches and improvement of the quality of large-scale educational assessments. *Information Security Journal: A Global Perspective, 33*(2), 1–15. https://doi.org/10.1080/19393555.2023.2284761
15. A. Nuraeni, Y. Nugraha, and M. E. Aminanto, "Revisiting Cyber Threats in Government Sectors: A Systematic Review of Attacks, Challenges, and Policy-Level Defenses," International Journal of Advances in Data and Information Systems, vol. 6, no. 2, pp. 447–459, Aug. 2025. doi: 10.59395/ijadis.v6i2.1404.
16. Bwiino, K., Mayoka, G. K., Nkamwesiga, L., Nyamadi, M., & Musenze, I. A. (2025). Information security behavior in higher education institutions. Journal of Information Security and Cybercrimes Research, 8(1), 01-16, DOI: 10.26735/YQBX3351
17. Omotunde, H., & Ahmed, M. (2023). A comprehensive review of security measures in database systems: Assessing authentication, access control, and beyond. Mesopotamian Journal of CyberSecurity, 2023, 115–133. https://doi.org/10.58496/MJCSC/2023/016
18. Dadheech, K., Choudhary, A., & Bhatia, G. (2018). De-militarized zone: A next level to network security. In Proceedings of the 2018 Second International Conference on Inventive Communication and Computational Technologies (ICICCT). https://doi.org/10.1109/ICICCT.2018.8473328
19. National Institute of Standards and Technology. (2023). Cybersecurity framework (Version 2.0). NIST. Available: https://www.nist.gov/cyberframework
20. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207
21. Scarfone, K., & Hoffman, P. (2009). Guidelines on firewalls and firewall policy (NIST Special Publication 800-41, Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-41r1
22. Whitman, M. E., & Mattord, H. J. (2022). Principles of information security (7th ed.). Cengage Learning.
23. Yang, J., Liang, L., & Qi, J. (2023). A practical non-intrusive cyber security vulnerability assessment method for cyber-insurance. In Proceedings of the 2023 IEEE International Conference on Data Science and Cybersecurity (DSC), pp. 261–269. doi: 10.1109/DSC59305.2023.00045
24. Hutchins, E. M., Cloppert, M. J., & Amin, R. M. (2011). Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. Leading Issues in Information Warfare & Security Research, 1(1), 80. Lockheed Martin. Available: https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf
25. Matherly, J. (2016). Complete guide to Shodan: Collect. Analyze. Visualize. Shodan LLC / Leanpub. Available: https://leanpub.com/shodan
26. Durumeric, Z., Adrian, D., Mirian, A., Bailey, M., & Halderman, J. A. (2015). A search engine backed by Internet-wide scanning. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15), pp. 542–553. ACM. https://doi.org/10.1145/2810103.2813703
27. Check Point Research. (2024). A closer look at Q3 2024: 75% surge in cyber attacks worldwide. Check Point Blog. Available: https://blog.checkpoint.com/research/a-closer-look-at-q3-2024-75-surge-in-cyber-attacks-worldwide/

28. Syed, N. F., Shah, S. W., Shaghaghi, A., Anwar, A., Baig, Z., & Doss, R. (2022). Zero Trust Architecture (ZTA): A comprehensive survey. IEEE Access, 10, 57143–57179. https://doi.org/10.1109/ACCESS.2022.3174679
29. Teerakanok, S., Uehara, T., & Inomata, A. (2021). Migrating to zero trust architecture: Reviews and challenges. Security and Communication Networks, 2021, Article 9947347. https://doi.org/10.1155/2021/9947347
30. Bitsight TRACE. (2023). Top cyber threats facing the education sector. Bitsight Research. Available: https://www.bitsight.com/blog/top-10-cyber-threats-facing-education-sector
31. Verizon. (2024). 2024 Data Breach Investigations Report (DBIR). Verizon Business. Available: https://www.verizon.com/business/resources/reports/dbir/