International Journal of Science, Technology, Engineering and Mathematics
Volume 6 Issue 3 September 2026
DOI: https://doi.org/10.53378/ijstem.353374

CVE-based security audit of open-source learning management systems in higher education using CVSS and OWASP

[1]Melissa T. Guillermo, [2]Eduardo R. Yu II & [3]Reagan B. Ricafort

Abstract

This study examined publicly disclosed Common Vulnerabilities and Exposures (CVE) records for Moodle, Chamilo, Canvas LMS, Open edX, and Sakai to assess their value for higher education security governance. A non-intrusive secondary-data audit analyzed 236 CVEs from 2018 to 3 June 2026 using platform, National Vulnerability Database (NVD) Published Date year, revised CVSS severity categories, and OWASP Top 10:2021 classifications. Six zero-score records were separated from low severity. A sensitivity analysis compared the full corpus with 211 NVD API-verified records. A blinded independent IT-faculty coder recoded 24 records; severity agreement was 100.0% (Cohen's kappa = 1.000), and OWASP agreement was 91.7% (Cohen's kappa = 0.814). Moodle and Chamilo represented 91.5% of the corpus, a disclosure-weighted pattern rather than a product ranking. A03 Injection (54.2%) and A01 Broken Access Control (31.4%) dominated. Among records with numeric CVSS scores, including six zero-score records, high/critical proportions were 45.3% (67/148) in the full corpus and 44.9% (66/147) in the NVD-verified subset. Canvas and Sakai findings were interpreted cautiously because of small denominators. Public CVE evidence is affected by disclosure asymmetry, incomplete CVSS metadata, NVD non-returned records, and deployment differences. Because identifiable CVSS version/vector metadata were not preserved, severity comparisons are descriptive rather than version-normalized. The audit offers a reproducible governance framework linking vulnerability patterns to patching, authorization review, content-ingestion controls, plugin governance, and supplier accountability without treating raw CVE counts as inherent security rankings.

Keywords: *cybersecurity governance, common vulnerabilities and exposures, cybersecurity, vulnerability assessment, OWASP, web application security*

Article History:

***Received**: June 9, 2026* ***Revised**: July 27, 2026*
***Accepted**: July 30, 2026* ***Published online**: August 10, 2026*

Suggested Citation:
Guillermo, M.T., Yu, E.R. & Ricafort, R.B. (2026). CVE-based security audit of open-source learning management systems in higher education using CVSS and OWASP. *International Journal of Science, Technology, Engineering and Mathematics*, 6(3), 1-26. https://doi.org/10.53378/ijstem.353374

About the authors:
[1]Corresponding author. Master of Science in Computer Science. Assistant Professor I, Dr. Filemon C. Aguilar Memorial College of Las Piñas, Las Piñas City, Philippines. Email: mrenrose@gmail.com
[2]Doctor of Information Technology, Licensed Professional Teacher, software engineer, and part-time faculty member. AMA University, Philippines. Email: eduardoryuii@gmail.com
[3]Doctor of Information Technology. Professor/faculty member, AMA University, Philippines. Email: reagan072001@gmail.com

1. Introduction

Learning management systems (LMS) support authentication, course enrollment, messaging, content distribution, file exchange, assessment, grading, and external integrations. Their centrality to teaching and administration means that a security failure can disrupt instruction while also exposing credentials, personal information, intellectual property, and institutional services. Continued reliance on hybrid and online provision has therefore made LMS security a governance issue rather than only a software-testing concern (Abdelkader et al., 2025; Apte & Sharma, 2021; Sadiqzade & Alisoy, 2025).

Higher education institutions face distinctive vulnerability-management constraints. University digital estates are heterogeneous, decentralized, and connected to identity providers, student information systems, research services, cloud applications, and third-party tools. Empirical work in the sector shows that cybersecurity response may also be weakened by fragmented threat-information sharing, management constraints, reputational concerns, and uneven coordination across institutions (Piazza et al., 2023). The 2024 EDUCAUSE Horizon Report likewise treats vulnerability management, data security, governance, and workforce capability as interdependent higher education priorities rather than isolated technical tasks (Robert et al., 2024). These conditions make publicly disclosed vulnerability evidence useful, but only when it is translated into operational decisions and interpreted with its limitations.

Recent incident evidence shows the operational consequences of LMS compromise. In May 2026, the U.S. Department of Education reported a cybersecurity incident affecting Canvas environments used by K-12 schools and institutions of higher education worldwide, involving unauthorized access to usernames, email addresses, course names, enrollment information, and messages. The Federal Bureau of Investigation separately reported that the attack interrupted service to educational institutions and students. These events are not treated as CVE-specific results in the present corpus; rather, they illustrate how LMS security failures can become institution-wide data-protection, service-continuity, phishing, vendor-management, and incident-communication problems (Federal Bureau of Investigation, 2026; Federal Student Aid, 2026).

Web-application weaknesses are especially relevant because LMS platforms process user-generated HTML, files, quizzes, messages, embedded media, authentication sessions, roles, and external requests. Prior studies have identified recurring risks involving injection, cross-site scripting (XSS), weak authentication, access-control errors, unsafe file handling, and

data protection (Aslan et al., 2023; Ben Fredj et al., 2020; Kasmawi et al., 2023; Mora Secaira et al., 2023; Nuhi et al., 2025). However, technical studies often focus on one installation, one product, or controlled penetration testing. Such work can identify concrete weaknesses but does not by itself show how publicly disclosed vulnerability classes are distributed across several LMS projects or how institutions should interpret unequal disclosure records.

CVE, CVSS, and OWASP provide complementary rather than interchangeable information. CVE gives a traceable public identifier and advisory record; CVSS supplies a severity-oriented triage signal when a usable score and version are available; and OWASP groups technical mechanisms into recognizable web-application control domains. Combining the three adds value because the same CVE count can be interpreted differently after severity completeness and weakness category are considered. The combination also makes it possible to connect an advisory record to institutional action: for example, a Critical or High score can inform remediation priority, while an A01 or A03 classification identifies the control family that requires review. The frameworks do not eliminate reporting bias, but their joint use produces a more decision-oriented audit than raw counts alone.

This study addresses three related gaps. First, it compares five open-source or open-core LMS platforms rather than limiting the analysis to a single deployment. Second, it integrates CVE records, available CVSS severity information, and OWASP Top 10:2021 coding in a governance-oriented audit. Third, it explicitly tests whether the main conclusions change when the 25 records not returned by the NVD API are excluded. The scientific contribution is therefore not a new vulnerability standard; it is a cross-platform, disclosure-aware synthesis that illustrates how public vulnerability metadata can be used without converting it into an unsupported product league table.

This study characterizes the distribution, timing, severity completeness, and OWASP categories of publicly disclosed vulnerabilities affecting Moodle, Chamilo, Canvas LMS, Open edX, and Sakai from 2018 to the 3 June 2026 dataset lock, and to translate the observed patterns into defensible higher education governance actions. Specifically, the study asked:

1. How are CVE records distributed across platforms and NVD published date years?
2. What severity patterns appear when scored, unscored, and zero-score records are distinguished?
3. Which OWASP categories and vulnerability mechanisms dominate?

4. Do the principal findings remain stable in the 211-record NVD-verified sensitivity corpus? and
5. What institutional actions are supported by the evidence?

2. Literature Review

2.1. LMS Cybersecurity and Institutional Vulnerability Governance

The LMS security research spans authentication controls, vulnerability scanning, penetration testing, software audits, and governance. Studies of Moodle and other educational platforms discuss authentication controls and report web-application risks, including XSS and injection, in specific LMS or university-web contexts (Banes et al., 2023; Kasmawi et al., 2023; Nuhi et al., 2025; Rossi et al., 2023). Because these studies are deployment- and method-specific, they do not estimate the cross-platform distribution of publicly disclosed weakness classes. Research on university websites and institutional LMS environments similarly reports recurring web-application risks, but the resulting evidence is usually local to the tested deployment and reflects its version, hosting model, plugins, authentication architecture, and configuration (Kasmawi et al., 2023; Mora Secaira et al., 2023).

Institutional governance studies broaden the problem. Higher education cybersecurity depends on standards, leadership, collaboration, asset ownership, incident communication, and the ability to coordinate remediation across technical and academic units (Agalit et al., 2023; Piazza et al., 2023; Robert et al., 2024; Shabbir et al., 2025; 2026). An LMS vulnerability is not resolved merely because a patch exists. Institutions must know which version and components are deployed, who owns the service, how roles and integrations are configured, and how quickly remediation can be tested and applied. Public CVE evidence is therefore most useful when it is embedded in an accountable vulnerability-management process.

Plugins, themes, connectors, and external tools complicate this process. They extend instructional functionality but also create additional code paths, maintenance dependencies, and disclosure channels. Moodle's own vulnerability-scanning guidance notes that plugin updates do not receive the same level of community scrutiny as core updates and advises administrators to verify whether findings concern installed plugins or core code (Moodle, 2026). Although plugin-only CVEs were excluded from the present core-platform corpus unless they belonged to the maintained official distribution, the plugin ecosystem remains a major deployment-level risk and is treated as a governance and future-research priority.

2.2. Comparison With Previous LMS Security Studies

The research gap becomes clearer when prior approaches are compared by scope, evidence source, and intended contribution. Table 1 shows that existing studies provide valuable platform-specific or multi-platform evidence, but few combine a five-platform public CVE corpus, severity completeness, OWASP coding, disclosure asymmetry, and an NVD-verified sensitivity analysis.

Table 1
Comparison of selected LMS cybersecurity studies and the present audit.

Study	Scope and data	Analytical approach	Main limitation for the present question	Distinct contribution of this study
Apte and Sharma (2021)	LMS cybersecurity concerns and participant roles	Conceptual and practice-oriented synthesis	Does not provide a cross-platform CVE distribution	Extends the governance concern with record-level cross-platform evidence
Akacha and Awad (2023)	Moodle, Chamilo, and ILIAS; 2018-2022 public vulnerabilities	Comparative vulnerability and stakeholder analysis	Narrower platform set and time window; no present NVD-verified sensitivity test	Adds five platforms, OWASP coding, June 2026 lock, metadata-completeness analysis, and sensitivity testing
Mora Secaira et al. (2023)	University-level LMS environment	Risk and vulnerability analysis	Single institutional context limits cross-platform inference	Uses public records across projects while explicitly retaining deployment limitations
Rossi et al. (2023)	Educational software in an experimental setting	Experimental vulnerability assessment	Technical testing is deployment- and method-specific	Complements testing evidence with disclosed-vulnerability governance signals
Abdelkader et al. (2025)	LMS audit framework in one university	Framework implementation and evaluation	Institution-specific evidence	Provides a reproducible comparative audit and governance mapping across five LMS projects
Present study	236 CVEs for Moodle, Chamilo, Canvas LMS, Open edX, and Sakai; 2018 to 3 June 2026	CVE + available CVSS + OWASP Top 10:2021; full and 211-record sensitivity corpora	Public disclosure evidence cannot measure all deployed risk	Shows which conclusions are robust to NVD non-returned records and links findings to higher education controls without ranking products

***Source*:** Apte and Sharma (2021), Akacha and Awad (2023), Mora Secaira et al. (2023), Rossi et al. (2023), and Abdelkader et al. (2025).

2.3. CVE, CVSS, and OWASP as Complementary Audit Frameworks

CVE and the NVD organize publicly disclosed vulnerability information, but the number of records is not a prevalence rate. CVSS provides a standardized base-score concept for triage, yet version changes and scoring inconsistency can affect comparison (Balsam et al., 2024; Forum of Incident Response and Security Teams, 2023; Wunder et al., 2024). OWASP Top 10:2021 offers a separate interpretive layer by grouping weaknesses into control-oriented categories. It is particularly relevant to LMS platforms because content input, file handling, authorization, authentication, and external requests are central platform functions (Open Web Application Security Project, 2021).

The frameworks are appropriate for this study because they operate at different levels. CVE is the unit and traceability mechanism; CVSS is a severity field used for operational triage; and OWASP is a weakness taxonomy used to interpret recurring control failures. The approach remains descriptive; it does not claim that the frameworks estimate exploit probability, institutional exposure, or the quality of a platform's development process. The outcomes depend on local versioning, exploit availability, configuration, compensating controls, and patch timing.

2.4. Disclosure Asymmetry and Analytical Framework

Public vulnerability evidence is shaped by researcher attention, project visibility, user base, self-hosted exposure, vendor reporting practices, coordinated disclosure, and database enrichment. A platform with more CVEs may have more weaknesses, but it may also have more transparent reporting and a larger research community. A platform with few records may be safer, less studied, more SaaS-oriented, or more likely to remediate privately. Missing scores and delayed NVD enrichment introduce additional asymmetry. Consequently, raw counts in this audit are interpreted as disclosure-weighted signals rather than normalized measures of inherent platform security.

The analytical framework links each CVE record to four primary dimensions: platform, NVD Published Date year, revised CVSS severity category, and OWASP Top 10:2021 category. Mechanism family provides additional details. These dimensions are then connected to governance actions: publication timing informs monitoring; severity and scoring completeness inform triage; A01 informs authorization review; A03 informs content and input controls; and disclosure asymmetry informs procurement and supplier-transparency

requirements. The framework's added value lies in making both the evidence and its uncertainty visible in the same audit.

3. Methodology

3.1. Research Design

The study used a quantitative descriptive-analytical design based on secondary data. No live scanning, penetration testing, exploitation attempt, or interaction with production LMS environments was conducted. The unit of analysis was a public CVE record linked to one of the five selected platforms. Descriptive counts and proportions were used because the public corpus is not a random sample of all vulnerabilities and cannot support causal or inferential claims about inherent product security.

The design was selected to support governance-oriented comparison rather than exploit development or platform ranking. It allows institutions to identify recurring disclosed weakness classes and metadata gaps while preserving a clear boundary between public advisory evidence and deployment-specific risk. Replication is supported through record-level identifiers, source URLs, search logic, CPE filters, coding rules, verification status, and analysis worksheets in Supplementary File S1.

3.2. Corpus and Platform Selection

The selected platforms were Moodle, Chamilo, Canvas LMS, Open edX, and Sakai. They were included because they are open-source or open-core LMS projects used in higher education, had at least one public CVE during the audit window, and had sufficient NVD, CVE Program, vendor, or community information for record verification. Mainly proprietary LMS products were not included because SaaS delivery and private disclosure practices are not directly comparable with an open-source CVE-based corpus. Third-party plugin-only issues were excluded unless the affected component belonged to the maintained core platform or official distribution.

The retained master file contained 236 unique post-search records. Candidate records had been collected and screened before the file was locked, but the raw pre-screening search log was not retained. Therefore, the earliest auditable count is 236, and the study does not claim that zero records were excluded before dataset lock or present the process as PRISMA flow. During the revision-stage recheck, no duplicates or additional off-target records were identified

in the locked file. This distinction resolves the apparent inconsistency between earlier narrative exclusions and the previous table showing 236 records at every retained stage.

3.3. Data Sources, Search Strategy, and Workflow

The CVE Program and the National Vulnerability Database (NVD) were the primary record sources (MITRE Corporation, n.d.; National Institute of Standards and Technology [NIST], n.d.); platform vendor or community advisories were used to verify scope and affected components. GitHub Security Advisories, CVE Details, and other third-party sources were used only for corroboration when a record was already traceable to an authoritative identifier or advisory. Search used platform names, CVE identifiers, CPE filters, and vendor/community channels. The dataset was locked on 3 June 2026, and NVD Published Date enrichment was retrieved at 2026-06-03T14:36:18Z. A record was retained when its title, description, CPE, advisory, or reference links supported relevance to the selected platform and scope.

Table 2

Platform-specific retrieval logic and primary verification sources

Platform	Representative search/CPE logic	Primary verification sources
Moodle	CVE ID lookups; (Moodle OR moodle/moodle) AND CVE; cpe:2.3:a:moodle:moodle:*	CVE Program; NVD; Moodle security advisories
Chamilo	CVE ID lookups; (Chamilo OR Chamilo LMS) AND CVE; cpe:2.3:a:chamilo:chamilo_lms:*	CVE Program; NVD; Chamilo announcements/community notices
Canvas LMS	CVE ID lookups; (Canvas LMS OR Instructure Canvas OR canvas-lms) AND CVE; cpe:2.3:a:instructure:canvas_lms:*	CVE Program; NVD; Instructure/Canvas advisories
Open edX	CVE ID lookups; (Open edX OR edx-platform OR openedx) AND CVE; openedx/edx-platform CPE and keyword variants	CVE Program; NVD; Open edX security notices/mailing lists
Sakai	CVE ID lookups; (Sakai LMS OR Apereo Sakai OR sakai) AND CVE; cpe:2.3:a:apereo:sakai:*	CVE Program; NVD; Sakai/Apereo security notices

***Source*:** Supplementary File S1, Search_Strategy sheet. Full URLs and record-level sources are retained in the workbook.

Figure 1

Locked-corpus audit and sensitivity-analysis workflow

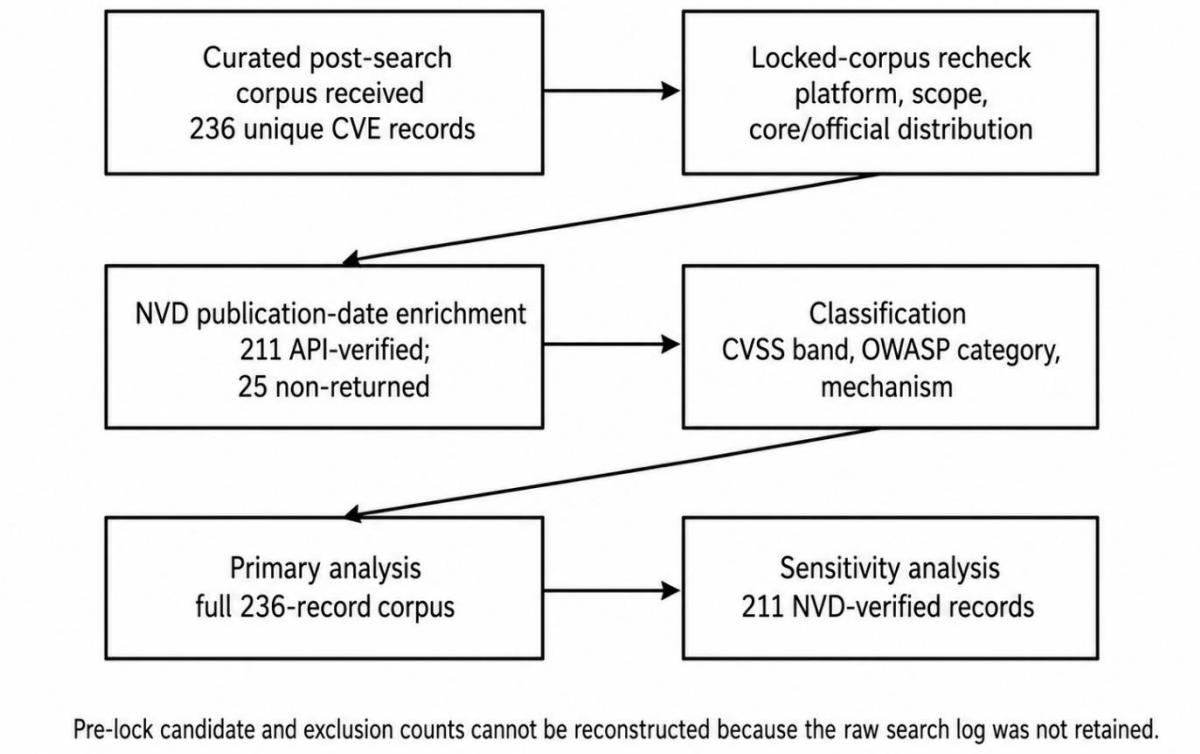

***Source*:** Generated from the locked-corpus workflow documented in Supplementary File S1.

Table 3

Auditable locked-corpus workflow and sensitivity corpus

Stage	Count	Operational interpretation
Curated post-search corpus received	236	Earliest count verifiable from retained files
Duplicate check within locked file	0 removed	All CVE IDs unique
Additional scope exclusions after lock	0 removed	No further off-target or out-of-scope records identified
Pre-lock exclusions	Not reconstructable	Raw candidate-search and exclusion log was not retained; no zero-exclusion claim is made
NVD publication-date verified	211	Returned by NVD CVE API 2.0 at the recorded retrieval time
NVD API non-returned	25	Retained with manual-verification flags; 24 Moodle and 1 Chamilo
Primary analysis	236	Full locked corpus
Sensitivity analysis	211	API-verified subset

***Source*:** Supplementary File S1, Curation_Workflow sheet

3.4. Coding Rules and Validity Controls

Each CVE was coded by platform, NVD Published Date year, CVSS severity category, OWASP Top 10:2021 category, and mechanism family. NVD Published Date was used when returned by the API. For 25 non-returned records, the previously curated year was retained with a manual-verification flag. The 2026 total is partial-year-to-date at the dataset lock.

Severity used the available numeric CVSS base score. Critical covered 9.0-10.0, High 7.0-8.9, Medium 4.0-6.9, and Low positive scores from 0.1-3.9. Six records with a numeric score of 0.0 were moved to a separate Zero score/unrated category rather than being automatically interpreted as Low. Records without a usable numeric score were coded Not Classified and treated as missing information. The locked file retained 148 numeric scores and 88 unscored records, but it did not preserve an identifiable CVSS metric version/vector for any record with a numeric CVSS score. A version-normalized comparison was therefore not possible and is not claimed.

OWASP categories followed the OWASP Top 10:2021 framework and CWE-to-category guidance where usable CWE data were present. XSS was assigned to A03 Injection. When the CWE and narrative were ambiguous, the category most directly supported by the described enabling mechanism was selected and documented; records without a defensible mapping were coded Unmapped. Multi-mechanism CVEs received one primary OWASP category based on the mechanism that directly enabled the principal reported impact, while secondary mechanisms were retained in the mechanism and rationale fields. This rule prevented double-counting one CVE across category totals.

Table 4

Revised coding rules for severity and ambiguous OWASP mappings

Variable	Code/rule	Operational definition
CVSS severity	Critical / High / Medium	9.0-10.0 / 7.0-8.9 / 4.0-6.9
CVSS severity	Low	Positive score from 0.1 to 3.9
CVSS severity	Zero score / unrated	Numeric score of 0.0; reported separately and not automatically treated as low risk
CVSS severity	Not Classified	No usable numeric score at dataset lock
OWASP	Ambiguous record	Use CWE and narrative; select only a defensible primary category or code Unmapped
OWASP	Multi-mechanism CVE	Assign one primary category by enabling mechanism; preserve secondary mechanisms in record-level fields

***Source*:** Supplementary File S1, Coding_Rubric and OWASP_Mapping_Notes sheets

Written rubric supported a blind independent recoding of 24 records (10.2% of the 236-record corpus) by an IT faculty member who was not involved in the primary classification. The coder worksheet displayed the source-based description, locked CVSS score, CVSS source note, CWE evidence, source URLs, and coding rubric, but did not display the primary severity or OWASP codes. The sample contained 12 Moodle, 10 Chamilo, one Open edX, and one Sakai record. Reliability was calculated from the initial pre-consensus classifications. Severity agreement was 24/24 (100.0%), with Cohen's kappa = 1.000. OWASP-category agreement was 22/24 (91.7%), with Cohen's kappa = 0.814. The two OWASP disagreements involved CVE-2023-30943 and CVE-2025-26532. No post-coding adjudication was documented, and the master-dataset codes were not changed after the reliability calculation. Record-level comparisons and formulas are provided in Supplementary File S1, Agreement and Reliability_Analysis sheets.

3.5. Data Analysis

The primary analysis described platform counts and shares, annual distribution based on NVD Published Date year, revised severity profiles, High/Critical proportions using total-record denominators and numeric-score denominators that include zero-score records, OWASP category distribution, and recurring mechanism families. Percentages for Canvas LMS and Sakai were treated as unstable descriptive values because their denominators were four and three records, respectively.

A sensitivity analysis repeated key comparisons after excluding the 25 records not returned by the NVD API. The 211-record verified corpus was compared with the full corpus on platform counts, High/Critical counts, High/Critical proportions among records with a numeric CVSS score (including zero-score records), and A01/A03 totals. This analysis assesses robustness to NVD date-verification status; it does not solve disclosure bias or missing CVSS-version metadata.

3.6. Research Ethics

Formal human-subject ethics approval was not applicable to this design because the study used public, non-personal vulnerability records and did not involve human participants, private records, live systems, active scanning, or exploitation. No exploit code or operational

attack instructions are reported. Source records remain traceable to public advisories, and the supplementary peer-review file was anonymized to preserve double-blind review.

4. Findings

4.1. Distribution Across Platforms

The full corpus contained 236 unique CVE records. Moodle accounted for 116 (49.2%) and Chamilo for 100 (42.4%), together representing 91.5% of the corpus. Open edX had 13 records, Canvas LMS four, and Sakai three. These are disclosure counts, not vulnerability-density rates.

Table 5

Distribution of CVE records across the five LMS platforms (n = 236).

Platform	CVE records	Share of corpus
Moodle	116	49.2%
Chamilo	100	42.4%
Open edX	13	5.5%
Canvas LMS	4	1.7%
Sakai	3	1.3%
Total	**236**	**100.0%**

***Source*:** Revised Supplementary File S1, Master_CVE_Dataset sheet

The Canvas and Sakai counts are too small for stable platform-level percentage comparisons. Their records are retained to preserve cross-platform coverage and to identify mechanisms that may still be operationally important, but no claim is made that their observed percentages represent the underlying platforms.

4.2. Temporal Distribution

Annual counts based on NVD Published Date year were lower from 2018 to 2022, rose in 2023, and remained visible in 2024 and 2025. Thirty-eight records had a 2026 NVD Published Date year, concentrated in Chamilo (28) and Moodle (10). Because the dataset was

locked on 3 June, the 2026 value is incomplete and cannot be compared directly with full calendar years. One Moodle record had an unspecified year.

Table 6

Annual CVE records by platform using NVD Published Date year (n = 236)

Year	Moodle	Chamilo	Canvas	Open edX	Sakai	Total
2018	3	4	0	0	0	7
2019	4	3	0	1	1	9
2020	3	0	2	4	0	9
2021	5	14	1	1	0	21
2022	5	9	0	1	0	15
2023	20	34	1	0	0	55
2024	27	8	0	4	1	40
2025	38	0	0	2	1	41
2026*	10	28	0	0	0	38
Unspecified	1	0	0	0	0	1
Total	**116**	**100**	**4**	**13**	**3**	**236**

*Partial-year-to-date at 3 June 2026

***Source*:** Revised Supplementary File S1, Annual_NVD_Date sheet

Figure 2

Annual CVE disclosures by platform using NVD Published Date year

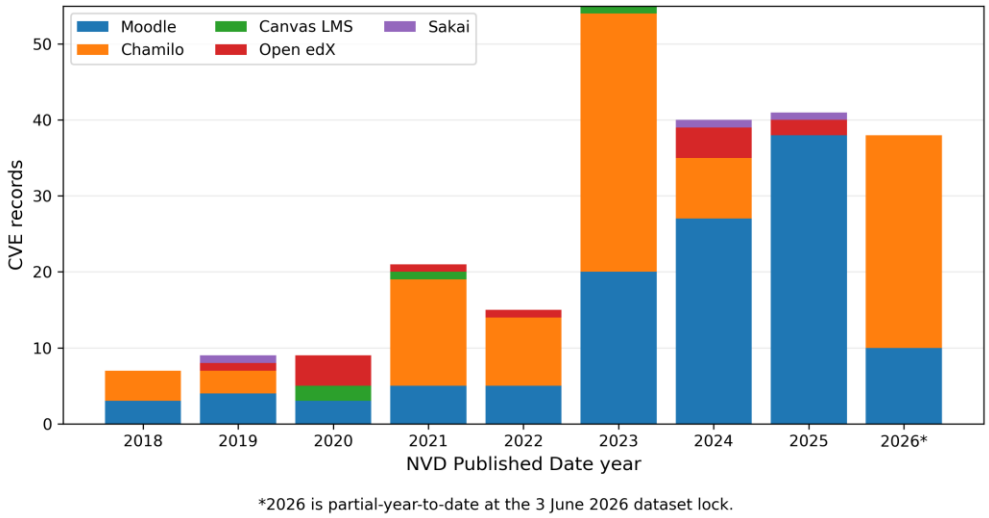

***Source*:** Generated from the locked revised dataset provided as Supplementary File S1

4.3. Revised CVSS Severity Profile

The revised classification contained 20 Critical, 47 High, 72 Medium, three Low, six Zero score / unrated, and 88 Not Classified records. Chamilo had the largest absolute number of High and Critical records. Moodle's total-record High/Critical proportion was 12.1%, but this value was strongly affected by 83 unscored records. Among the 33 Moodle records with a numeric CVSS score, 14 were High or Critical (42.4%).

Table 7

Revised CVSS severity profile by platform (n = 236). H+C% uses total platform records.

Platform	Crit.	High	Med.	Low	Zero	Not class.	Total	H+C%
Moodle	2	12	19	0	0	83	116	12.1%
Chamilo	17	28	42	3	6	4	100	45.0%
Canvas LMS	0	1	2	0	0	1	4	25.0%
Open edX	1	5	7	0	0	0	13	46.2%
Sakai	0	1	2	0	0	0	3	33.3%
Total	**20**	**47**	**72**	**3**	**6**	**88**	**236**	**28.4%**

Zero = numeric score 0.0 reported separately. Not class. = no usable numeric score. Canvas and Sakai percentages have very small denominators.
***Source*:** Revised Supplementary File S1, Revised_Severity_Summary sheet

Figure 3

Revised CVSS severity profile by platform

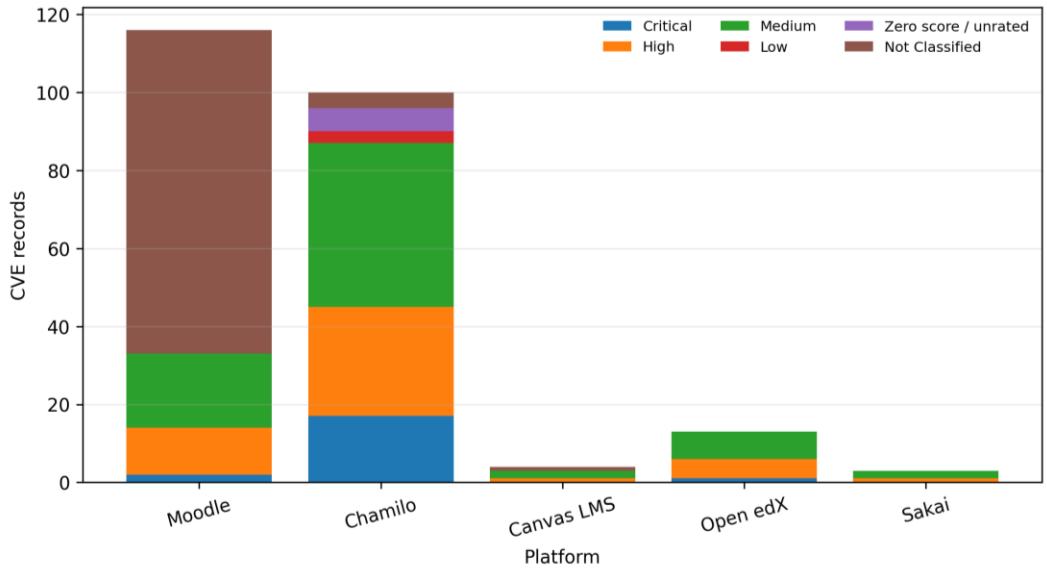

***Source*:** Generated from the locked revised dataset provided as Supplementary File S1

4.4. NVD-Verified Sensitivity Analysis

The 211-record sensitivity corpus produced nearly the same severity pattern among records with a numeric CVSS score as the full corpus. High/Critical records decreased from 67 to 66, and the High/Critical proportion among numeric-score records, including zero-score records, changed from 45.3% (67/148) to 44.9% (66/147). A03 remained the largest OWASP category and A01 remained second. Differences were concentrated in Moodle because 24 of the 25 NVD API non-returned records were Moodle records.

Table 8

Sensitivity comparison between the full corpus and NVD API-verified records

Measure	Full corpus	Verified corpus	Interpretation
Records	236	211	25 non-returned records excluded
High + Critical	67	66	Absolute difference of one record
Records with numeric CVSS scores (zero scores included)	148	147	One excluded record had a numeric CVSS score
H+C among numeric-score records (zero scores included)	45.3% (67/148)	44.9% (66/147)	Substantively stable under the same denominator definition
A03 Injection	128	119	Largest category in both corpora
A01 Broken Access Control	74	62	Second-largest category in both corpora
Moodle records	116	92	Accounts for 24 of the 25 exclusions
Chamilo records	100	99	Accounts for one exclusion

***Source*:** Revised Supplementary File S1, Sensitivity_Analysis sheet.

4.5. CVSS Metadata Completeness

A comparable-version severity analysis could not be performed from the locked dataset. Although 148 records had numeric base scores, all 148 had an unspecified CVSS metric version/vector field in the retained file; 88 records had no usable score. Consequently, the severity bands describe triage categories based on available numbers and should not be interpreted as a normalized comparison of CVSS v2.0, v3.x, or v4.0 scores.

Table 9

CVSS metadata completeness in the locked corpus

Metadata condition	Records	Analytical consequence
Usable numeric base score	148	Included in the numeric-score denominator; zero-score records are included
No usable numeric base score	88	Coded Not Classified
Identifiable CVSS version/vector preserved	0	No version-normalized analysis possible
Numeric score but version unspecified	148	Severity bands used only as descriptive triage indicators

***Source*:** Revised Supplementary File S1, CVSS_Version_Limitations sheet

4.6. OWASP Categories and Mechanism Families

A03 Injection covered 128 records (54.2%), followed by A01 Broken Access Control with 74 (31.4%). A07 Identification and Authentication Failures accounted for 11; A08 Software and Data Integrity Failures and A10 Server-Side Request Forgery had five each; and 13 were Unmapped. A03 was especially visible in Chamilo, whereas A01 was the largest Moodle category.

Table 10

OWASP Top 10:2021 categories by platform (n = 236)

Category	Moodle	Chamilo	Canvas	Open edX	Sakai	Total
A01 Broken Access Control	53	15	1	4	1	74
A03 Injection	43	75	1	8	1	128
A07 Identification and Authentication Failures	10	1	0	0	0	11
A08 Software and Data Integrity Failures	0	4	0	0	1	5
A10 Server-Side Request Forgery	1	2	2	0	0	5
Unmapped	9	3	0	1	0	13
Total	116	100	4	13	3	236

***Source*:** Revised Supplementary File S1, Master_CVE_Dataset sheet

Figure 4

Overall OWASP Top 10:2021 category distribution

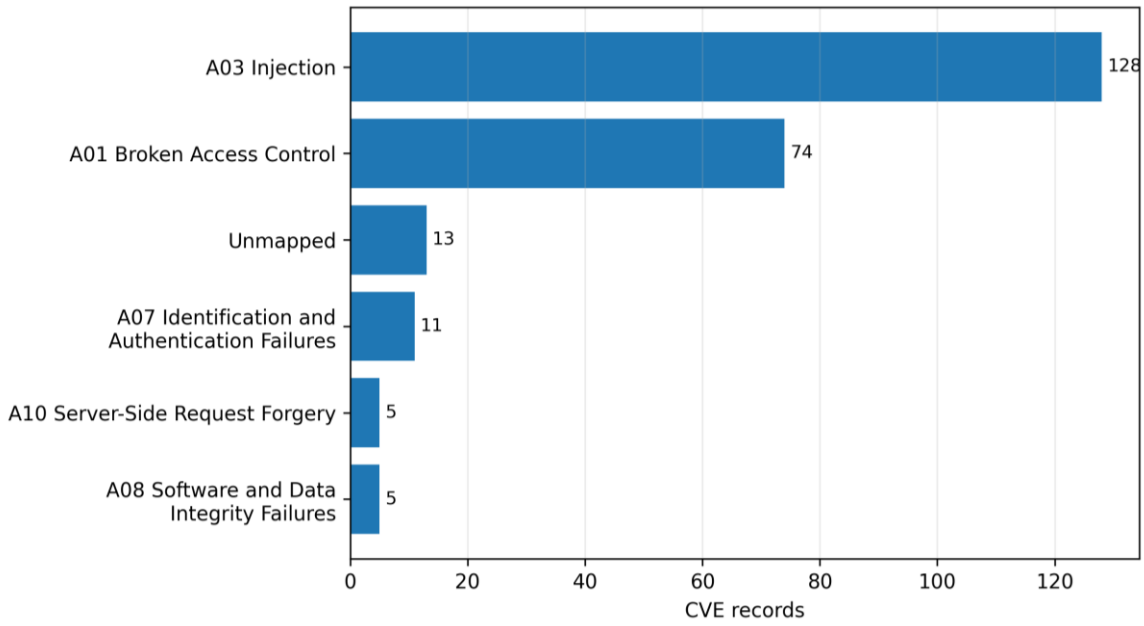

***Source*:** Generated from the locked revised dataset provided as Supplementary File S1

At the mechanism level, XSS represented 73 records (30.9%), authorization bypass/IDOR 48 (20.3%), and RCE or unsafe upload leading to execution 36 (15.3%). Lower-frequency mechanisms included SQL injection, SSRF, insecure deserialization, command/code injection, and XXE. Canvas had only four records, but two involved SSRF, illustrating why a small count should not be equated with negligible impact.

5. Discussion

5.1. Scientific Contribution and Relationship to Previous Research

The study's principal scientific contribution is a disclosure-aware cross-platform synthesis rather than a claim that combining established frameworks is itself novel. Earlier work has demonstrated LMS authentication weaknesses, local web vulnerabilities, and comparative public vulnerability patterns (Akacha & Awad, 2023; Banes et al., 2023; Mora Secaira et al., 2023; Nuhi et al., 2025). The present analysis extends that literature by applying one record-level coding structure to five projects, updating the window to June 2026, separating missing and zero severity information, and testing the main conclusions against the NVD-verified subset. This makes the contribution methodological and interpretive: it shows

how to use public vulnerability records as governance evidence while exposing the conditions that can distort comparison.

The sensitivity analysis is consistent with the robustness of the numeric-score comparison. Excluding 25 non-returned records changed the High/Critical proportion among records with a numeric CVSS score, including zero-score records, by less than one percentage point and did not change the ordering of the two dominant OWASP categories. Thus, the conclusions about severity among numeric-score records and the importance of A03/A01 are not artifacts of retaining the non-returned records. However, the reduction from 116 to 92 Moodle records confirms that platform shares remain sensitive to database-return and enrichment patterns. Robustness in category dominance does not convert raw platform totals into normalized security rankings.

5.2. Disclosure Asymmetry and Temporal Interpretation

Moodle and Chamilo dominated the corpus, but several explanations remain plausible: larger self-hosted footprints, more visible advisory channels, stronger research attention, coordinated disclosure practices, or genuine differences in the number of weaknesses discovered. Public CVE volume therefore has at least two meanings, exposure signal and transparency signal. The finding is useful for monitoring because it identifies where public advisories are concentrated, but it cannot establish a denominator such as installed instances, code size, supported versions, or researcher effort.

The post-2022 rise should be interpreted similarly. It may reflect increased discovery or exposure, but it may also reflect publication timing, CNA workflow, delayed NVD enrichment, or grouped advisory releases. The concentration of the partial 2026 records in two platforms reinforces this caution. A stronger longitudinal study would need reporter, CNA, advisory, patch-release, and affected-version timelines before attributing the pattern to worsening code quality.

5.3. Severity Completeness and the Meaning of Uncertainty

The Moodle result illustrates why missing severity data can alter the risk narrative. Its High/Critical proportion was 12.1% when all 116 records were used but 42.4% among its 33 records with a numeric CVSS score. Treating 83 Not Classified records as if they were low risk would suppress the apparent severity. For institutional teams, an unscored CVE should

trigger local assessment rather than automatic de-prioritization: affected versions, internet exposure, exploit preconditions, known exploitation, business criticality, and compensating controls may be more informative than an absent score.

Separating the six zero-score records also improves interpretive discipline. A numeric zero can reflect unusual scoring or record status and does not necessarily provide the same evidence as a positive Low score. The revised table prevents those records from being used as proof of low impact. At the same time, the lack of preserved CVSS versions limits cross-record comparability. Prior research shows that CVSS versions and scoring judgments can produce inconsistent severity readings (Balsam et al., 2024; Wunder et al., 2024). The audit therefore uses CVSS as a triage aid, not an objective product-quality measure.

5.4. Why A03 Injection and A01 Broken Access Control Matter for LMS Operations

The dominance of A03 is consistent with the functional design of LMS platforms. They accept rich text, files, questions, messages, imports, embedded media, and externally sourced content. These features increase opportunities for unsafe rendering, inadequate validation, injection chains, and execution through uploaded or restored content. The practical implication is broader than filtering one form field. Institutions need layered content-ingestion controls, supported versions, secure defaults, file-type and MIME validation, safer rendering, content security policy where feasible, attachment scanning, and rapid review of components that process imported content.

A01 reflects the equally complex authorization environment of higher education. LMS platforms distinguish administrators, course creators, instructors, teaching assistants, graders, students, guests, external tools, and delegated support personnel. Permission boundaries may change through course copying, enrollment automation, delegated administration, role overrides, and integrations. Regular role and permission review is therefore necessary even when the core platform is patched. The high Moodle A01 count is not treated as a Moodle-only design conclusion; it is evidence that authorization deserves sustained governance attention across deployments.

Small-platform findings require a different level of inference. Two of four Canvas records involved SSRF, but the denominator is too small to estimate a stable platform tendency. The proper conclusion is that institutions using Canvas should not ignore SSRF-related controls merely because the total public record count is small. The same caution applies

to Sakai. Mechanism-level relevance can be operationally meaningful even when platform percentages are statistically fragile.

5.5. Evidence-Linked Governance Implications

Table 11 links each recommendation to an empirical result and an authoritative control source. This structure avoids generic recommendations by showing why each action follows from the audit.

Table 11

Governance actions linked to empirical findings and security standards

Empirical finding	Governance action	Operational evidence or standard
67 High/Critical records; 45.3% of records with a numeric CVSS score (zero scores included)	Define emergency, expedited, and routine remediation targets; verify installation and compensating controls	NIST SP 800-40 Rev. 4 (Souppaya & Scarfone, 2022); NIST CSF 2.0 (NIST, 2024)
88 Not Classified records and no preserved CVSS versions	Require local triage for unscored CVEs and retain complete metric/vector metadata in the vulnerability register	NIST CSF 2.0 risk assessment and continuous monitoring (NIST, 2024)
A03 Injection = 54.2%; XSS and RCE/upload mechanisms prominent	Harden content ingestion, rendering, file validation, upload/restore processes, and application security testing	OWASP Top 10:2021 (Open Web Application Security Project, 2021); NIST SSDF 1.1 (Scarfone et al., 2022)
A01 Broken Access Control = 31.4%	Review role definitions, overrides, delegated administration, enrollment rules, APIs, and external-tool permissions	OWASP Top 10:2021 (Open Web Application Security Project, 2021); ISO/IEC 27001:2022 and ISO/IEC 27002:2022 (International Organization for Standardization, 2022a, 2022b)

Empirical finding	Governance action	Operational evidence or standard
Disclosure concentrated in Moodle and Chamilo; unequal public visibility	Monitor vendor/community advisories and require transparent vulnerability and incident communication in procurement	NIST SSDF supplier communication (Scarfone et al., 2022); institutional supplier governance
Plugin ecosystems create deployment-specific exposure	Maintain a versioned inventory of plugins/themes/integrations; remove unsupported extensions and monitor their advisories	Moodle vulnerability-scanning guidance (Moodle, 2026); NIST CSF 2.0 asset management (NIST, 2024)
Canvas and Sakai have very small denominators	Use record-level mechanism review rather than platform percentage ranking	Study limitation and risk-based local assessment
25 NVD API non-returned records	Cross-check NVD, CVE Program, and vendor sources; flag unresolved metadata rather than silently filling it	Revised supplementary verification workflow

***Source*:** Empirical results from this study; standards and guidance cited in the References.

In practice, institutions should maintain an LMS asset register that includes core version, hosting model, plugins, themes, integrations, authentication mechanisms, exposed endpoints, service owner, patch status, and supplier contacts. Critical and High advisories should be reviewed against local exposure rather than applied mechanically. Procurement and hosted-service agreements should require timely vulnerability disclosure, patch notification, incident communication, advisory transparency, and defined service-level commitments. NIST patch-management guidance frames patching as enterprise preventive maintenance, while the SSDF provides a shared vocabulary for communicating secure-development and supplier expectations (Scarfone et al., 2022; Souppaya & Scarfone, 2022).

5.6. Limitations

First, the corpus includes only public CVE records and excludes privately discovered or privately remediated issues. Second, disclosure practices differ across self-hosted, community, and SaaS-oriented products, so raw counts cannot normalize platform security. Third, the pre-lock candidate-search and exclusion log was not retained; the workflow is

auditable from the locked 236-record file onward but cannot reconstruct earlier attrition. Fourth, 25 records were not returned by the NVD API at retrieval, although the sensitivity analysis indicates that the main severity and OWASP conclusions are stable. Fifth, exact CVSS versions and vectors were not preserved, preventing a version-normalized comparison.

Sixth, public advisory text requires judgment. The independent reliability exercise was limited to 24 records and did not include a Canvas LMS record. Although severity agreement was perfect and OWASP agreement was strong, the two OWASP disagreements show that ambiguous source narratives can still affect category assignment. The reported coefficients are pre-consensus because no post-coding adjudication was documented. Seventh, plugin-only records were generally excluded even though plugins, themes, and integrations can dominate real deployment risk. Eighth, the design does not include exploitability, known-exploitation status, affected installation counts, patch latency, local configuration, or compensating controls. Ninth, the partial 2026 period and small Canvas/Sakai denominators limit temporal and percentage inference. The results are descriptive and non-causal.

5.7. Future Research

Future studies should preserve the full candidate-search log and use a reproducible screening protocol with explicit pre-lock exclusion reasons. Complete NVD CVSS metric objects should be re-fetched so that v3.x-only or version-specific sensitivity analyses can be conducted. Reliability testing should be extended to a larger prespecified stratified sample that includes every platform, followed by a documented adjudication procedure while retaining the pre-consensus coefficients. Reporter, CNA, advisory, patch-release, and affected-version dates should be integrated to distinguish genuine discovery trends from disclosure batches and enrichment delays.

The platform scope should also expand to plugin and theme ecosystems, official connectors, and external learning tools. Public CVE audits could be combined with ethics-approved controlled testing of representative deployments, software bills of materials, configuration reviews, patch-latency data, and institutional incident records. Comparative work on proprietary or hosted LMS products should use vendor trust portals, security bulletins, procurement evidence, and incident-notification commitments to reduce the asymmetry created by private remediation. Longitudinal monitoring is needed to determine whether the post-2022 pattern persists after full-year 2026 data become available.

6. Conclusion

Scientifically, this study contributes a five-platform, 236-record synthesis that integrates public CVE traceability, available CVSS triage, OWASP weakness categories, and explicit disclosure-asymmetry interpretation. It confirms that A03 Injection and A01 Broken Access Control dominate the disclosed corpus, while avoiding the unsupported conclusion that platforms with more records are inherently less secure.

Methodologically, the revision illustrates the importance of separating unscored and zero-score records, reporting CVSS-version incompleteness, and testing findings against a verified-record sensitivity corpus. The near-identical High/Critical proportions among records with a numeric CVSS score, including zero-score records, in the full and verified corpora support the stability of the main severity conclusion, but the concentration of non-returned records in Moodle confirms that platform totals remain disclosure dependent. The blinded 24-record second-coder exercise adds chance-corrected evidence of coding reproducibility, with Cohen's kappa values of 1.000 for severity and 0.814 for OWASP categories.

Practically, higher education decision-makers can use the findings to prioritize patch governance, role and permission review, content-ingestion controls, plugin lifecycle management, complete vulnerability metadata, and supplier transparency. The audit is best used as an operational monitoring framework and procurement input, not as a product scorecard. Larger all-platform reliability samples, complete CVSS-vector enrichment, plugin analysis, and deployment-level evidence are necessary to strengthen comparative inference.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was not supported by any funding.

AI Declaration

The authors declare the use of Artificial Intelligence (AI) in preparing this paper. In particular, the authors used ChatGPT for grammar checking, language clarity, low-intervention paraphrasing, formatting review, and consistency checking of manuscript sections. The tool was not used to fabricate data, generate the CVE dataset, create unsupported findings, or replace the authors' interpretation of the results. Python was also used as a research-support and data-processing tool for organizing the CVE dataset, checking record counts, computing descriptive statistics and percentages, preparing tables, and generating figures used in the manuscript. The authors take full responsibility for reviewing, verifying, and editing all AI-assisted and software-supported outputs.

ORCID

Melissa T. Guillermo – https://orcid.org/0009-0009-4051-8393

Eduardo R. Yu II – https://orcid.org/0009-0004-5050-5805

Reagan B. Ricafort– https://orcid.org/0009-0002-0816-1522

References

Abdelkader, M. A., Mazen, S. A., & Helal, I. M. A. (2025). Implementation and evaluation of a proposed framework for auditing learning management systems in an Egyptian university. *International Journal of Computer Information Systems and Industrial Management Applications, 17*, 746-764. https://doi.org/10.70917/ijcisim-2025-0045

Agalit, M. A., Chakir, E. M., Issam, T., & Khamlichi, Y. I. (2023). A review of cybersecurity management standards applied in higher education institutions. *International Journal of Safety and Security Engineering, 13*(6), 1109-1116. https://doi.org/10.18280/ijsse.130614

Akacha, S. A., & Awad, A. I. (2023). Enhancing security and sustainability of e-learning software systems: A comprehensive vulnerability analysis and recommendations for stakeholders. *Sustainability, 15*(19), 14132. https://doi.org/10.3390/su151914132

Apte, T., & Sharma, S. (2021). Progression in cyber security concerns for learning management systems: Analyzing the role of participants. In V. Bali, V. Bhatnagar, D. Aggarwal, S. Bali, & M. J. Sousa (Eds.), *Cyber-physical, IoT, and autonomous systems in Industry 4.0*. CRC Press. https://doi.org/10.1201/9781003146711

Aslan, O., Aktug, S. S., Okay, M. O., Yilmaz, A. A., & Akin, E. (2023). A comprehensive review of cyber security vulnerabilities, threats, attacks, and solutions. *Electronics, 12*(6), 1333. https://doi.org/10.3390/electronics12061333

Balsam, A., Nowak, M., Walkowski, M., Oko, J., & Sujecki, S. (2024). Comprehensive comparison between versions CVSS v2.0, CVSS v3.x and CVSS v4.0 as vulnerability severity measures. In *2024 24th International Conference on Transparent Optical Networks (ICTON)*. IEEE. https://doi.org/10.1109/ICTON62926.2024.10647452

Banes, V., Ravariu, C., Appasani, B., & Srinivasulu, A. (2023). A novel two-factor authentication scheme for increased security in accessing the Moodle e-learning platform. *Applied Sciences, 13*(17), 9675. https://doi.org/10.3390/app13179675

Ben Fredj, O., Cheikhrouhou, O., Krichen, M., Hamam, H., & Derhab, A. (2020). An OWASP Top Ten driven survey on web application protection methods. In J. Garcia-Alfaro, J. Leneutre, N. Cuppens, & R. Yaich (Eds.), *Risks and security of Internet and systems* (pp. 235-252). Springer. https://doi.org/10.1007/978-3-030-68887-5_14

Federal Bureau of Investigation. (2026, May 15). *ShinyHunters: Cyber criminal group attacks learning management system* (Alert No. I-051526-PSA). Internet Crime Complaint Center. https://www.ic3.gov/PSA/2026/PSA260515

Federal Student Aid. (2026, May 12). *Technology security alert—Ongoing cybersecurity incident involving the Canvas Learning Management System* (GENERAL-26-27). U.S. Department of Education. https://fsapartners.ed.gov/knowledge-center/library/electronic-announcements/2026-05-12/technology-security-alert-ongoing-cybersecurity-incident-involving-canvas-learning-management-system-updated-may-29-2026

Forum of Incident Response and Security Teams. (2023). *Common Vulnerability Scoring System version 4.0: Specification document*. https://www.first.org/cvss/v4.0/specification-document

International Organization for Standardization. (2022a). *Information security, cybersecurity and privacy protection - Information security management systems - Requirements (ISO/IEC 27001:2022)*. https://www.iso.org/standard/27001

International Organization for Standardization. (2022b). *Information security, cybersecurity and privacy protection - Information security controls (ISO/IEC 27002:2022)*. https://www.iso.org/standard/75652.html

Kasmawi, Hidayasari, N., & Mansur. (2023). Vulnerability analysis using OWASP ZAP on higher education websites. *AIP Conference Proceedings, 2665*(1), 030015. https://doi.org/10.1063/5.0153145

MITRE Corporation. (n.d.). *CVE Program*. https://www.cve.org/

Moodle. (2026). *Vulnerability scans*. https://docs.moodle.org/502/en/Vulnerability_scans

Mora Secaira, J., Díaz Ocampo, R., & Samaniego Mena, E. (2023). Analysis of risks and vulnerabilities in a university-level LMS system. *Migration Letters, 20*(S8), 1252-1262. https://www.migrationletters.com/index.php/ml/article/view/5090

National Institute of Standards and Technology. (n.d.). *National Vulnerability Database*. https://nvd.nist.gov/

National Institute of Standards and Technology. (2024). *The NIST Cybersecurity Framework (CSF) 2.0 (NIST CSWP 29)*. https://doi.org/10.6028/NIST.CSWP.29

Nuhi, A., Ajruli, N., Idrizi, F., Imeri, F., & Memeti, A. (2025). Securing MVC-based LMS platforms: Addressing authentication, XSS, and injection vulnerabilities. *Journal of Natural Sciences and Mathematics of UT, 10*(19-20), 312-321. https://doi.org/10.62792/ut.jnsm.v10.i19-20.p3198

Open Web Application Security Project. (2021). *OWASP Top 10:2021 - The ten most critical web application security risks*. https://owasp.org/Top10/

Piazza, A., Vasudevan, S., & Carr, M. (2023). Cybersecurity in UK universities: Mapping (or managing) threat intelligence sharing within the higher education sector. *Journal of Cybersecurity, 9*(1), tyad019. https://doi.org/10.1093/cybsec/tyad019

Robert, J., Muscanell, N., Arbino, N., McCormack, M., & Reeves, J. (2024). *2024 EDUCAUSE Horizon Report: Cybersecurity and Privacy Edition*. EDUCAUSE. https://library.educause.edu/resources/2024/9/2024-educause-horizon-report-cybersecurity-and-privacy-edition

Rossi, D., Bressan, L., Campos, F., Oliveira, A., & Stroele, V. (2023). Educational software and security vulnerabilities: An experimental study. In *Anais do XXXIV Simposio Brasileiro de Informatica na Educacao* (pp. 264-276). SBC. https://doi.org/10.5753/sbie.2023.234860

Sadiqzade, Z., & Alisoy, H. (2025). Cybersecurity and online education: Risks and solutions. *Luminis Applied Science and Engineering, 2*(1), 4-12. https://doi.org/10.69760/lumin.20250001001

Scarfone, K., Souppaya, M., & Dodson, D. (2022). *Secure Software Development Framework (SSDF) version 1.1: Recommendations for mitigating the risk of software vulnerabilities (NIST SP 800-218)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-218

Shabbir, A., Deraman, A., Hassan, M. N. B., Sarker, K. U., & Kamal, S. (2025). Prioritizing non-functional requirements and influencing factors for API quality framework: An industry approach. *International Journal of Advanced Computer Science and Applications, 16*(9). https://doi.org/10.14569/IJACSA.2025.0160934

Shabbir, A., Deraman, A., Hassan, M. N. B., Sarker, K. U., & Kamal, S. (2026). Enhancing API quality: A comprehensive review of non-functional requirements for quality-centric framework development. *KSII Transactions on Internet and Information Systems, 20*(2), 853–875. https://doi.org/10.3837/tiis.2026.02.011

Souppaya, M., & Scarfone, K. (2022). *Guide to enterprise patch management planning: Preventive maintenance for technology (NIST SP 800-40 Rev. 4)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-40r4

Wunder, J., Kurtz, A., Eichenmuller, C., Gassmann, F., & Benenson, Z. (2024). Shedding light on CVSS scoring inconsistencies: A user-centric study on evaluating widespread security vulnerabilities. In *2024 IEEE Symposium on Security and Privacy* (pp. 1102-1121). IEEE. https://doi.org/10.1109/SP54263.2024.00058