
Annual Penetration Testing Review 2025



Contents

1. Introduction	03
2. Executive Summary	04
3. Methodology	06
4. Key Findings	10
5. Assessment Types Comparison	18
6. Industry Benchmarks	27
7. Industries Deep-Dive	33
– E-commerce and Retail	
– Finance and Fintech	
– Tech, Software and SaaS	
– Healthcare and Pharmaceutical	
– Oil, Gas and Energy	
8. Compliance Insights	38
– SOC 2 Penetration Tests	
– ISO 27001 Penetration Tests	
– PCI DSS Penetration Tests	
9. Conclusion	44

1.Introduction

This report presents a comprehensive analysis of penetration testing engagements conducted by Blaze Information Security throughout 2025.

Unlike theoretical vulnerability databases or automated scanning results, the findings presented here represent confirmed security weaknesses discovered through manual penetration testing by experienced security consultants. Each vulnerability was validated, exploited where appropriate, and documented with sufficient detail to enable effective remediation.

For security leaders, this report offers crucial benchmarking data to contextualize their organization's security posture against industry peers. For technical teams, it provides actionable intelligence on the most prevalent attack vectors and vulnerability patterns. For compliance and risk management professionals, it delivers evidence-based insights into how regulatory frameworks translate into real-world security outcomes.



2. Executive Summary

This report analyzes the penetration testing engagements conducted by Blaze Information Security between January and mid-November 2025, covering 145 organizations from 11 industries. Across these assessments, each vulnerability represents a weakness that was actively discovered, confirmed, and exploited when possible, during a real-world penetration testing engagement.

**660**Total Pentests
Performed**3,294**Total Vulnerabilities
Identified**145**

Unique Clients

Timeline



Jan 2025

Nov 2025

Several clear patterns emerge from this year's data.

First, application-layer weaknesses continue to dominate the vulnerability landscape.

Web applications, APIs, and mobile clients account for the vast majority of findings and for the most severe risks observed. While this reflects the scope of testing performed, it also mirrors where modern attack surfaces have concentrated: internet-facing services, multi-tenant platforms, and API-driven architectures.


Second, the most dangerous vulnerabilities are not the most technically complex.

Critical and high-severity findings account for a minority of total issues, but they are disproportionately driven by a small set of recurring weaknesses. Access control failures, authorization bypasses, insecure design decisions, and protection mechanism failures consistently enable unauthorized access, privilege escalation, and sensitive data exposure. These are not edge-case bugs, but systemic failures in how trust boundaries and permissions are defined and enforced.

Third, most exploitable vulnerabilities are reachable remotely and require little to no user interaction.

Over 85% of vulnerabilities in the dataset are exploitable over a network, and more than half require no prior authentication. Even when authentication is required, exploitation frequently involves legitimate users abusing overly broad permissions rather than attackers bypassing perimeter defenses. These findings reinforce a recurring theme throughout the report: identity exists, but authorization is frequently incomplete, inconsistent, or incorrect.

Remote exploitability of findings



At the same time, the data shows incremental progress. Many organizations have implemented baseline security controls, particularly around perimeter exposure, encryption and authentication. Fully unauthenticated, “open door” vulnerabilities are less common than in previous years. However, this progress is offset by persistent weaknesses inside trusted boundaries, where controls exist but fail under real-world conditions.

Compliance-driven assessments (SOC 2, ISO 27001, PCI DSS) reinforce this conclusion. Environments scoped for regulatory requirements rarely fail due to missing controls, but frequently due to incorrect enforcement of those controls in application logic, session handling, authorization checks, and configuration management. In practice, compliance improves baseline hygiene, but does not prevent design-level and logic-level security failures.

Across industries, maturity varies, but no sector is immune. Regulated industries tend to exhibit fewer low-effort external exposures, yet often show a higher concentration of high-impact findings related to identity, authorization, and privilege management. Technology-driven sectors, while more security-aware, continue to struggle with scale-induced complexity and inconsistent enforcement across rapidly evolving systems.

Taken together, the 2025 findings point to a clear conclusion: **the primary security challenge is no longer building controls, but ensuring they are correct, consistent, and resilient across real application workflows.** The most effective risk reduction does not come from adding more tools or controls, but from strengthening authorization models, minimizing implicit trust and enforcing defense-in-depth where it matters most.

The remainder of this report explores these patterns in detail, providing benchmarks, industry comparisons and practical insights to help organizations understand where they stand – and where their defenses are most likely to fail.

3. Methodology

Assessment Types Included

All data included in this report was collected between January and mid-November 2025. The analysis covers penetration testing engagements executed during this period and reflects the state of real-world security findings observed throughout most of the year.

Timeline and Scope


The dataset encompasses a wide range of offensive security engagements, covering multiple attack surfaces and testing methodologies. The cybersecurity assessments we performed most often are Web Application Security Testing (60.80%), API Security Testing (15.5%) and Mobile Application Security (8.8%), followed by Infrastructure Security Testing (Internal and External), Red Teaming, and Cloud Security Assessments. We also conducted 21 specialized assessments, including source code reviews, Wi-Fi security testing, IoT penetration testing and LLM security assessments.

Around 70% of our assessments are followed by a retest to confirm if remediating measures were successfully implemented. However, the findings from those projects were not included in this report.

Client Sample Characteristics


The significant concentration in the Tech and Software & SaaS sectors enables deep analytical insights into the security maturity of high-growth technology companies, digital-native organizations, and cloud-based service providers. Simultaneously, the dataset provides comparative benchmarking data for e-Commerce sector and regulated industries such as Finance, Healthcare, and Energy, where security requirements are often more prescriptive.

Clients by Industry



Clients range from small and medium-sized organizations to large enterprises, operating in the United States, Europe, Asia-Pacific and LATAM regions.

Clients by Size



Data Collection and Analysis

All vulnerability data was collected through **VulnKeep**, Blaze's Pentest as a Service (PTaaS) platform. VulnKeep is a platform that provides our clients with instant visibility into validated findings, audit-ready reporting, and integrated workflows.

For analytical purposes, the data was exported and structured into a consolidated dataset (Excel-based), enabling statistical analysis and cross-correlation across multiple dimensions, including project type, CWE, Vulnerability Code, Title, Severity, Likelihood and Impact, CVSS code, CVSS Vector and others.

Limitations and Potential Biases

While this report is based on a substantial real-world dataset, a few limitations and potential sources of bias should be considered when interpreting the findings:

- Web application and API security assessments account for the majority of tests performed. As a result, vulnerabilities related to these attack surfaces are likely overrepresented compared to areas such as internal infrastructure, cloud, or red team engagements.
- The findings reflect what clients chose to test. Assets that were out of scope, unknown, or intentionally excluded are not represented, which may understate exposure in certain environments.
- Technology, Software, and SaaS organizations are disproportionately represented in the dataset. As a result, industry-wide conclusions may not fully generalize to sectors with fewer samples.

Despite these limitations, the dataset provides a robust and realistic view of the vulnerabilities and attack patterns observed during penetration testing engagements in 2025.



4. Key Findings

1. Vulnerability Trends

The analysis of **3,294 confirmed vulnerabilities** identified across **660 penetration tests** in 2025 reveals a vulnerability landscape strongly shaped by application-layer weaknesses, access control failures, and systemic misconfigurations.

While the majority of findings fall into lower and medium severity categories, a recurring subset of high-impact vulnerabilities continues to pose significant risk.

Vulnerabilities by Severity

From a severity perspective, the dataset is largely composed of **Low and Medium severity findings**, which together account for the majority of reported issues (56.5%).

However, **84 Critical** and **358 High severity vulnerabilities** were also identified, representing approximately **13.5%** of all findings. They are frequently associated with authorization flaws, broken security controls, and design-level weaknesses that enable unauthorized access, privilege escalation, or large-scale data exposure.

How We Categorize Severities

Critical: CVSS 9.0 – 10.0 (Fix ASAP)

High: CVSS 7.0 – 8.9 (Stop and evaluate)

Medium: CVSS 4.0 – 6.9 (Evaluate in context)


Low: CVSS 0.1 – 3.9 (Minor issues, hard to exploit)

None/Informational: CVSS 0.0 (No impact)

Vulnerability Criteria Classification

Severity	Description
CRITICAL	Results in the compromise of the system and the data it handles. Can be exploited by an unskilled attacker using publicly available tools and exploits. Requires immediate remediation.
HIGH	Likely to result in the compromise of the system and the data it handles.
MEDIUM	Does not lead to the immediate compromise of the system, but when chained with other issues can bring serious security risks. Nevertheless, it is advisable to fix them accordingly.
LOW	Does not pose an immediate risk, and even when chained with other vulnerabilities, it is less likely to cause a serious impact.
INFO	Does not pose an immediate risk, but requires continuous surveillance so that it doesn't become a liability through ill-use or future modifications to the system.

Vulnerabilities Distribution by Severity



This distribution reinforces a recurring pattern observed in penetration testing: the most dangerous vulnerabilities are not always the most frequent, but rather those that undermine fundamental trust and control boundaries within systems.


Focusing specifically on **Critical and High severity findings**, access control failures emerge as the most significant driver of risk. **CWE-284 (Improper Access Control)** stands out as the most prevalent vulnerability category among high-impact findings.

While 164 instances of CWE-284 were identified overall, **64 of these were classified as Critical or High severity**, making it the single most influential weakness in terms of real-world exploit potential.

The prominence of access control failures among severe findings underscores a recurring challenge: authorization logic is often complex, distributed, and tightly coupled to business rules, making it especially prone to subtle but dangerous implementation flaws.

Most Common Vulnerabilities Identified

Across all engagements, Blaze identified **206 distinct CWE categories**, reflecting a wide diversity of weaknesses. These were the top 10 most frequently observed vulnerabilities:



- CWE-200: Exposure of Sensitive Information to an Unauthorized Actor

CWE-799: Improper Control of Interaction Frequency

CWE-284: Improper Access Control

CWE-693: Protection Mechanism Failure

CWE-209: Generation of Error Message Containing Sensitive Information

CWE-204: Observable Response Discrepancy
- CWE-20: Improper Input Validation

CWE-307: Improper Restriction of Excessive Authentication Attempts

CWE-327: Use of a Broken or Risky Cryptographic Algorithm

CWE-79: Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting')

Taken together, these findings point to systemic issues in how applications handle sensitive data, enforce authorization decisions, and implement defensive controls. Rather than isolated coding mistakes, many of the most common vulnerabilities reflect inconsistent security assumptions, missing guardrails, or security controls that are present but incorrectly enforced.

Most of our findings were categorized according to the OWASP Top 10 by the pentesters. This categorization is not linear, since sometimes a CWE can fit into multiple OWASP categories. This is our Top 5 OWASP vulnerabilities:

Rank	OWASP Category	Findings	%
1	A5 – Security Misconfiguration	681	20.67%
2	A7 – Identification & Authentication Failures	328	9.96%
3	A4 – Insecure Design	317	9.62%
4	A1 – Broken Access Control	284	8.62%
5	A3 – Injection	239	7.26%

Types of Pentest with the Most Vulnerabilities

When analyzing vulnerability distribution by assessment type, web application testing clearly accounts for the **largest absolute number of findings**. Web applications alone contributed nearly **60% of all identified vulnerabilities**, followed by mobile applications and APIs.

However, raw volume does not tell the full story, since Web Application was also the type of assessment we performed the most, by far. When findings are normalized by the number of assessments performed, certain engagement types exhibit a significantly higher **vulnerability density**. In particular, **cloud security assessments**, despite their smaller sample size, revealed the highest average number of vulnerabilities per project: 14.40, followed by Mobile application and Red team engagements.

Overall, the average number of vulnerabilities per project was 4.99 with the median being four.

Most Common Vulnerabilities

Web Applications



CWE-200

Exposure of Sensitive Information

Mobile Applications



CWE-295 - Improper Certificate Validation



CWE-250 - Execution with Unnecessary Privileges

API



CWE-20

Improper Input Validation

Cloud



CWE-693

Protection Mechanism Failure

In **web applications**, information exposure remains the dominant issue, often stemming from overly verbose responses, misconfigured access controls, or unintended data leakage through application logic.

In **mobile applications**, CWE-295 and CWE-250 are almost equally as frequent.


For **API assessments**, the most prevalent recurring issue is **improper input validation**, frequently manifesting as missing constraints, unsafe defaults, or insufficient validation of request parameters. These weaknesses often enable abuse scenarios such as excessive resource consumption, data exposure, or unauthorized object access.

In **cloud environments**, recurring findings are primarily linked to **protection mechanism failures**, reflecting misconfigured security controls, overly permissive identity and access management policies, or missing enforcement of defense-in-depth principles.

2. Attack Vector Analysis

To better understand how the 2025 findings translate into real-world exploitability, we analyzed the CVSS Vector String of all vulnerabilities found.

Base Score Metrics Distribution



Remote Exploitability Dominates

The clearest signal in the dataset is how frequently vulnerabilities are reachable remotely: **85.1%** of scored findings are exploitable over a network rather than requiring local or physical access.

The result aligns with what we typically see in modern environments, where the most exposed attack surfaces are public-facing applications, APIs,

and internet-accessible services, all of which naturally map to “network” in CVSS terms (i.e., reachable over at least one network hop). It also matches established threat models for initial access: MITRE explicitly notes that adversaries frequently gain footholds by exploiting weaknesses in public-facing servers/applications.

Preconditions for Exploitation

We also found that:

Attack Complexity is nearly evenly split between **Low (49.8%)** and **High (50.2%)**. This suggests that many findings are not “single-request compromises,” but also that a large fraction are still reliably exploitable without special conditions.

Privileges Required are None for 53.0% of vulnerabilities, meaning exploitation often does not require an authenticated foothold. This is an important indicator of exposure, particularly in externally reachable systems.

User Interaction is None in the majority of cases (83.5%). In other words, exploitation can typically be performed directly by an attacker, without the intervenience of an user.

Together, these metrics indicate that a meaningful portion of the dataset represents vulnerabilities that are both remotely reachable and operationally practical to exploit.

The Impact of Vulnerabilities

70.5% of vulnerabilities affect confidentiality to some degree (Low or High). This far exceeds Integrity (35.6% affected) and availability (26.9% affected).

This distribution is consistent with an application-heavy testing portfolio, where the most damaging outcomes frequently involve:

- Unauthorized access to sensitive customer data
- Exposure of personally identifiable information (PII)
- Disclosure of intellectual property or trade secrets
- Leakage of authentication credentials or API keys

In contrast, pure availability disruption (denial of service) represents a smaller portion of findings, likely because:

- Application-layer DoS vulnerabilities are increasingly mitigated by cloud infrastructure and CDN protections
- Penetration testing typically focuses on unauthorized access rather than service disruption scenarios
- Many rate-limiting and resource exhaustion issues are classified as Medium severity rather than High availability impact

High-impact vulnerability analysis:

A total of 692 vulnerabilities score “High” in at least one impact dimension (**confidentiality, integrity, or availability**), representing 21.0% of all findings. Of these: 364 vulnerabilities score “High” in at least two of the three dimensions (10.9% of total)

These multi-dimensional high-impact findings represent the most severe threats, as they enable attackers to simultaneously exfiltrate data, modify systems, and potentially disrupt operations

These are the Top 5 vulnerabilities with at least 2 High Impact Scores:

CWE	Name	Count (≥2 High Impact Scores)
284	Improper Access Control	39
79	Improper Neutralization of Input During Web Page Generation (Cross-site Scripting)	36
250	Execution with Unnecessary Privileges	30
269	Improper Privilege Management	26
295	Improper Certificate Validation	19




5. Assessment Types Comparison

Comparing assessment types provides important insights into where vulnerabilities concentrate, how severe they tend to be, and how different testing approaches uncover risk.

Web Application Security Testing accounts for 60.9% of all engagements in our 2025 dataset. Next are API Security Testing (15.5%) and Mobile Application Security (8.8%), followed by Infrastructure Security Testing (Internal and External), Red Teaming, and Cloud Security Assessments. In addition, 21 specialized assessments were conducted, including source code reviews, Wi-Fi security testing, IoT penetration testing and LLM security assessments.

Distribution of Security Testing Types (2025)



Severity Distribution by Assessment Type

Assessment Type	Critical	High	Medium	Low	Info
Web Application Security Testing	2.7%	11.9%	30.2%	43.0%	12.3%
Mobile & Client-Side Security Testing	2.1%	6.3%	42.1%	42.1%	7.3%
Mixed / Specialized Assessment	2.3%	12.4%	27.1%	45.9%	12.4%
API Security Testing	0.8%	3.1%	15.4%	57.9%	22.8%
Infrastructure Penetration Testing	3.1%	14.9%	34.2%	38.5%	9.3%
Cloud Security	1.7%	1.7%	26.7%	68.3%	1.7%
Red Teaming	14.3%	38.1%	28.6%	7.1%	11.9%
Source Code Review	0.0%	0.0%	9.1%	72.7%	18.2%

Across assessment types, the proportion of High and Critical findings varies markedly depending on scope and attacker model. Red Team engagements predictably exhibit the highest concentration of severe issues, with High and Critical findings forming a substantial share of all discoveries, reflecting the focus on exploit chains and business-impact scenarios rather than isolated weaknesses.

Infrastructure penetration testing also shows an elevated High and Medium mix, as network exposure, trust relationships, and privilege boundaries frequently enable escalation paths.


Web application security testing shows the widest severity spread, including the highest absolute number of Critical (53) and High (237) vulnerabilities. Web and mobile application assessments are dominated by Medium-severity findings, complemented by a consistent High-severity baseline, indicating widespread design, access control, and authentication weaknesses that are impactful but not always immediately exploitable.

In contrast, API and cloud security assessments show comparatively low proportions of High and Critical issues, with most findings falling into Medium or below, driven largely by configuration and hardening gaps rather than direct compromise vectors.

Common Findings by Assessment Type


The **web results** skew toward information exposure and access control problems, plus patterns that point to rate-limiting/anti-automation gaps. In practice, this usually means sensitive data can leak through endpoints, error messages, or overly broad responses, and users can reach functions or records they shouldn't. Session lifecycle and authentication issues reinforce that account security and authorization boundaries are frequent weak points. Overall, the chart suggests that "classic" web risks still dominate: protecting data returned by the app, tightening authorization checks, and hardening auth/session handling tend to yield the biggest risk reduction.

Web




API findings concentrate around authorization and data exposure: endpoints returning more data than necessary, weak object-level authorization, and authentication gaps. This typically reflects services that assume trusted callers, rely on client-supplied identifiers, or lack consistent access control enforcement across endpoints. Error handling shows up too, which often enables information leakage useful for enumeration. The practical takeaway is that APIs frequently fail at "who can do what to which object," and at minimizing what they reveal. Stronger authorization patterns, input validation, and consistent error responses usually have an outsized impact on API security.

API




Mobile results highlight confidentiality and crypto hygiene issues: sensitive data exposure, missing/weak encryption, insecure local storage, and insecure transport/certificate validation patterns. That combination usually indicates risks around data at rest on the device and data in transit (e.g., tokens, PII, session artifacts). The presence of caching/logging storage issues suggests that apps may unintentionally persist sensitive data in places accessible to other apps, backups, or device compromise scenarios. The overall message is that mobile security often hinges on protecting secrets and user data across the full lifecycle: storage, logging, networking, and authentication flows.

Mobile




Infrastructure findings point to identity and credential weaknesses: weak password requirements, authentication/access control issues, and credential handling problems. This typically maps to insecure defaults, credential reuse, insufficient hardening of privileged accounts, and gaps in segmentation or service configuration. Privilege management and authorization issues imply pathways for lateral movement or escalation once a foothold exists. In practice, these results usually mean that improving credential policy, access boundaries, privileged access management, and service hardening can materially reduce the likelihood of compromise and the blast radius of a successful intrusion.

Infrastructure




Red team results are dominated by credential and access control themes: weak passwords, authentication issues, privilege management, and credential exposure. This is consistent with adversary-style testing: attackers often win through identity weaknesses rather than “exotic” exploits. The table suggests that organizational defenses are most stress-tested around identity controls, privilege boundaries, and authorization enforcement. The implication is that investments in strong authentication, resistant credential policies, privileged access controls, and detection around credential abuse and lateral movement are likely to address the most common real-world compromise paths observed in these exercises.

Red team




Cloud findings emphasize permissioning and data exposure: incorrect permissions, access control failures, privilege management, and protection mechanism failures. This pattern typically reflects misconfigured IAM, overly broad roles, mis-scoped policies, and security controls not consistently enforced across cloud resources. Because cloud environments are highly permission-driven, minor configuration errors can lead to outsized exposure. The practical meaning is that cloud security improvements often come from tightening identity/permission models, enforcing least privilege, validating resource policies, and standardizing guardrails (e.g., policy-as-code, continuous configuration monitoring) to prevent misconfigurations from drifting into production.

Cloud (AWS -> Cloud)



Code review findings skew toward data exposure and access control, with occasional hard-coded credentials and cryptography pitfalls. This usually indicates systemic design/implementation issues: sensitive data handling decisions embedded in code paths, authorization logic that isn't centralized, and secrets creeping into repositories or configs. Even when counts are smaller, these are high-leverage because they point to defects that can repeat across multiple services or releases. The implication is that secure coding practices, secret management, and reusable authorization patterns (plus code scanning and review checklists) can prevent entire classes of issues before they reach deployed systems.

Code Review




MITRE Top 25 vs Our Top Critical Vulnerabilities

MITRE has unveiled its latest [Top 25 Most Dangerous Software Weaknesses](#), behind over 39,000 security vulnerabilities disclosed between June 2024 and June 2025.

The list was developed in collaboration with the Homeland Security Systems Engineering and Development Institute (HSSDI) and the Cybersecurity and Infrastructure Security Agency (CISA), which jointly manage and sponsor the Common Weakness Enumeration (CWE) program.

To create the ranking, MITRE scored each weakness based on its severity and frequency after analyzing 39,080 CVE Records.

Between January and November 2025, we found 17 of these vulnerabilities in our projects:



In total, we found 894 vulnerabilities from MITRE's Top 25, including:

- 307 findings of CWE-200 Exposure of Sensitive Information
- 164 findings of CWE-284 Improper Access Control
- 113 findings of CWE-20 Improper Input Validation
- 72 findings of CWE-79 Cross-site Scripting
- 60 findings of CWE-434 Unrestricted Upload of File with Dangerous Type

Despite some similarities with the MITRE Top 25, our Top 25 vulnerabilities look a little different:


In total, six weakness categories appear in both MITRE's Top 25 and Blaze's Top 25 rankings. This overlap demonstrates that several of the most operationally significant weaknesses in penetration testing are also recognized in publicly disclosed CVE datasets.

MITRE's rankings are strongly influenced by memory safety and low-level implementation flaws, such as out-of-bounds writes and use-after-free conditions, which frequently generate CVE disclosures and affect system-level software.

By contrast, Blaze's dataset is dominated by application-layer weaknesses: authorization failures, data exposure, protection mechanism failures, rate-limiting issues, and logic-driven flaws. These weaknesses are typically context-dependent and emerge during manual analysis of application workflows, tenant separation models and access control enforcement.

MITRE reflects weaknesses that become publicly visible across products and platforms. Penetration testing reveals weaknesses that are exploitable within specific system implementations, even if they never result in CVE disclosure.

Taken together, the comparison shows that while foundational weakness categories remain consistent, the operational reality of modern web and cloud environments shifts the center of gravity toward authorization, data handling, and application logic as primary drivers of exploitable risk.




6. Industry Benchmarks

Industry Distribution by Clients and Projects

The industry distribution of our 2025 client sample shows a clear concentration in Tech, Software & SaaS, followed by Finance & FinTech, which together represent the majority of organizations assessed.


Distribution of Clients by Industry

N = 145



However, when looking at the total number of penetration testing projects per industry, a different pattern emerges. E-Commerce & Retail (242 projects) and Finance & Fintech (182 projects) account for the highest volumes of engagements, surpassing Tech, Software & SaaS in total project count, as organizations in these sectors tend to commission multiple assessments per year, often across different applications, platforms, or business units.

Distribution of Projects by Industry



A key factor behind this divergence is organization size. Many of the largest organizations in our dataset operate in the E-Commerce & Retail sector, where complex digital ecosystems, high transaction volumes, and frequent releases drive a higher testing cadence. Large retail environments often include multiple brands, regional platforms, mobile applications, APIs, and third-party integrations, each requiring separate or recurring security validation.

Assessment Type by Industry

Industry	Total projects	Web Application Security Testing	API Security Testing	Mobile Application Security Testing	Infrastructure Security Testing	Red Team Assessments	AWS Cloud Security	Mixed/ Specialized/ Others
e-Comm & Retail	242	50.41%	31.4%	7.85%	0.83%	0.0%	0.0%	7.85%
Finance & Fintech	182	58.79%	11.54%	13.74%	3.85%	1.65%	0.55%	9.89%
Healthcare & Pharmaceutical	29	65.52%	0.0%	3.45%	6.9%	3.45%	3.45%	17.24%
Education	5	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Oil, Gas & Energy	17	70.59%	5.88%	5.88%	11.76%	5.88%	0.0%	0.0%
Transport & Mobility	16	62.5%	0.0%	25.0%	6.25%	0.0%	0.0%	6.25%
Telecoms	19	47.37%	0.0%	5.26%	42.11%	0.0%	0.0%	5.26%
Tech, Software & SaaS	117	62.39%	3.42%	5.98%	5.98%	0.85%	1.71%	19.66%
Public / Other Services	13	84.62%	0.02%	0.0%	15.38%	0.0%	0.0%	0.0%
Manufacturing	4	50.0%	0.0%	0.0%	25.0%	0.0%	0.0%	25.0%
Insurance	8	62.5%	0.0%	0.0%	25.0%	0.0%	0.0%	12.5%

Across nearly all industries, Web Application Security Testing is the dominant assessment type, typically accounting for 50–85% of projects.

The table also highlights clear sector-specific testing profiles. e-Commerce & Retail shows the most diversified portfolio, with a large share of API testing (31.4%) alongside web testing, consistent with integration-heavy platforms and transaction-driven systems. Finance & Fintech has the highest relative concentration of mobile (13.7%) and a meaningful share of API (11.5%) testing, reflecting mobile-first user journeys and service-oriented architectures.


In contrast, Telecoms stand out for its unusually high share of Infrastructure Security Testing (42.1%), suggesting greater emphasis on network exposure, segmentation, and internal attack paths compared to other sectors. Healthcare & Pharmaceutical and Tech, Software & SaaS include the largest “Mixed/Specialized/Others” components (17.2% and 19.7% respectively), indicating broader scopes and a higher frequency of specialized engagements beyond standard web/API testing.

Severity Distribution by Industry

Severity distributions vary significantly across industries, reflecting differences in threat exposure, regulatory pressure, system complexity, and security maturity.

Industries such as Oil, Gas & Energy, Manufacturing, Education, and Public / Other Services exhibit a higher proportion of Critical and High severity findings, despite representing a smaller share of total projects. It suggests that lower testing frequency does not necessarily correlate with lower risk, and that concentrated exposure or legacy systems may amplify impact when vulnerabilities are present

In contrast, e-Commerce & Retail and Finance & Fintech show lower proportions of Critical and High severity findings, but contribute a large absolute number of vulnerabilities due to their overall testing volume and attack surface size.



Industries with the Highest % of Critical/high Findings

It’s also interesting to notice that some industries have a significantly higher percentage of critical and high severity findings, such as Education, Public services and insurance companies, for instance.

Industry	High + Critical %
Education	32.0%
Public / Other Services	29.2%
Insurance	27.6%
Manufacturing	26.1%
Oil, Gas & Energy	25.0%
Healthcare & Pharmaceutical	18.4%
Telecoms	16.3%
Tech, Software & SaaS	14.3%
Finance & Fintech	11.9%
Transport & Mobility	9.6%
e-Comm & Retail	9.1%

Vulnerabilities by Sector

As expected, the sectors where we found the most vulnerabilities are those where we performed the most tests. However, we can still see that some industries have a higher number of vulnerabilities per project, such as Insurance, by far, Oil, Gas & Energy and Tech, Software & SaaS.

Industry	% of all findings	Projects (total)	Avg. vulnerabilities per project
e-Comm & Retail	31.12%	242	4.24
Finance & Fintech	25.53%	187	4.50
Tech, Software & SaaS	22.92%	119	6.34
Healthcare & Pharmaceutical	5.43%	30	5.97
Oil, Gas & Energy	3.40%	17	6.59
Transport & Mobility	2.85%	16	5.88
Telecoms	2.79%	19	4.84
Insurance	2.31%	8	9.50
Public / Other Services	2.19%	13	5.54
Education	0.76%	5	5.00
Manufacturing	0.70%	4	5.75

Most Common Vulnerabilities by Industry


Finally, we can see that the most common vulnerabilities vary slightly across sectors. Even so, CWE-200 is the most common in almost all sectors, except Manufacturing, where this vulnerability does not even appear in the top 3.

Industry	Vulnerabilities (total)	Top 1	Top 2	Top 3
e-Comm & Retail	1025	CWE-200 (91, 8.9%)	CWE-693 (78, 7.6%)	CWE-20 (58, 5.7%)
Finance & Fintech	841	CWE-200 (67, 8.0%)	CWE-799 (66, 7.8%)	CWE-284 (40, 4.8%)
Tech, Software & SaaS	755	CWE-200 (65, 8.6%)	CWE-284 (39, 5.2%)	CWE-79 (38, 5.0%)
Healthcare & Pharmaceutical	179	CWE-200 (18, 10.1%)	CWE-284 (13, 7.3%)	CWE-307 (6, 3.4%)
Oil, Gas & Energy	112	CWE-200 (19, 17.0%)	CWE-284 (12, 10.7%)	CWE-799 (10, 8.9%)
Transport & Mobility	94	CWE-200 (7, 7.4%)	CWE-250 (7, 7.4%)	CWE-295 (6, 6.4%)
Telecoms	92	CWE-200 (14, 15.2%)	CWE-307 (5, 5.4%)	CWE-693 (5, 5.4%)
Insurance	76	CWE-200 (15, 19.7%)	CWE-284 (8, 10.5%)	CWE-20 (4, 5.3%)
Public / Other Services	72	CWE-200 (8, 11.1%)	CWE-434 (7, 9.7%)	CWE-204 (4, 5.6%)
Education	25	CWE-200 (2, 8.0%)	CWE-204 (2, 8.0%)	CWE-284 (2, 8.0%)
Manufacturing	23	CWE-284 (4, 17.4%)	CWE-1390 (2, 8.7%)	CWE-1391 (2, 8.7%)

7. Industries Deep-Dive

E-commerce and Retail


1. Vulnerabilities severity distribution



2. Average number of vulnerabilities per project

4.24

3. Top 10 vulnerabilities




4. Attack vector insights

Findings in e-commerce and retail are predominantly network-exploitable, with nearly four out of five vulnerabilities reachable remotely. A notable secondary share of adjacent and local vectors reflects exposure in internal admin panels, partner integrations, and backend systems supporting commerce workflows. Most vulnerabilities require no prior authentication and no user interaction, indicating that issues often reside in publicly accessible components rather than user-driven flows. Impact is skewed toward confidentiality and integrity loss, consistent with weaknesses in access control, session handling, and configuration of customer-facing systems. Overall, the attack surface is shaped by high exposure and complex integration ecosystems rather than physical or insider-only access.

Finance and Fintech


1. Vulnerabilities severity distribution



2. Average number of vulnerabilities per project

4.50

3. Top 10 vulnerabilities




4. Attack vector insights

Finance and fintech findings are overwhelmingly network-based, but this sector shows a higher proportion of local attack vectors than most others, reflecting the presence of internal systems, privileged administrative interfaces, and operational tooling. Vulnerabilities frequently require no user interaction, but a meaningful subset assumes some level of authentication, aligning with flaws discovered post-login rather than purely external attacks. Compared to other industries, finance shows higher confidentiality and integrity impact, consistent with the nature of financial data and transaction logic. Scope change remains rare, indicating that most findings affect individual components rather than cascading across trust boundaries.

Tech, Software and SaaS


1. Vulnerabilities severity distribution



2. Average number of vulnerabilities per project

6.34

3. Top 10 vulnerabilities




4. Attack vector insights

The SaaS sector exhibits the highest concentration of network-exploitable vulnerabilities, with almost all findings reachable remotely. Attack complexity is evenly split, but most issues require no user interaction and minimal privileges, highlighting weaknesses in APIs, application logic, and default configurations rather than user-dependent flows. Compared to regulated industries, SaaS shows a slightly higher rate of scope change, reflecting microservice architectures and shared trust boundaries. Impact is dominated by confidentiality and integrity rather than availability, reinforcing that risks primarily affect data exposure and authorization rather than service disruption.

Healthcare and Pharmaceutical


1. Vulnerabilities severity distribution



2. Average number of vulnerabilities per project

5.97

3. Top 10 vulnerabilities




4. Attack vector insights

Healthcare and pharmaceutical findings remain largely network-driven, but this sector stands out for its higher presence of local and physical attack vectors. This aligns with hybrid environments combining web systems, on-prem infrastructure, and specialized internal applications. A larger share of vulnerabilities requires some privileges, indicating risks emerging after initial access rather than exclusively at the perimeter. Impact distribution is more balanced across confidentiality, integrity, and availability, reflecting the operational sensitivity of clinical and research systems. The data shows a broader attack surface shaped by mixed deployment models rather than purely internet-facing assets.

Oil, Gas and Energy


1. Vulnerabilities severity distribution



2. Average number of vulnerabilities per project

6.59

3. Top 10 vulnerabilities



4. Attack vector insights


Despite operating in critical infrastructure contexts, the majority of findings in oil, gas, and energy are still remotely exploitable over the network, typically through management interfaces, exposed services, or cloud-connected systems. User interaction is rarely required, and most issues assume low or no privileges, indicating weaknesses in access controls and configuration rather than insider-only flaws. Compared to other industries, this sector shows a higher availability impact, reflecting operational systems where service disruption carries tangible risk. Physical and adjacent vectors exist but remain marginal, suggesting that discovered vulnerabilities are concentrated in IT-connected layers rather than isolated OT environments.

8. Compliance Insights

SOC 2 Penetration Tests

SOC 2 pentests are heavily skewed toward externally facing attack surfaces, especially web applications, with infrastructure tests typically used to validate broader Trust Services Criteria (network access, segmentation, baseline security). Types of assessments driven by SOC 2 compliance needs include web application pentests, web and API pentests, external (cloud) pentests, infrastructure pentests, and red teaming.

SOC 2 Findings by Severity:



Top 10 vulnerabilities in SOC 2 pentest

CWEN	Name	% of findings (SOC 2 set)
CWE-284	Improper Access Control	12.9%
CWE-200	Exposure of Sensitive Information to an Unauthorized Actor	8.6%
CWE-209	Generation of Error Message Containing Sensitive Information	7.1%
CWE-203	Observable Discrepancy	4.3%
CWE-327	Use of a Broken or Risky Cryptographic Algorithm	4.3%
CWE-345	Insufficient Verification of Data Authenticity	4.3%
CWE-601	URL Redirection to Untrusted Site (Open Redirect)	4.3%
CWE-613	Insufficient Session Expiration	4.3%
CWE-639	Authorization Bypass Through User-Controlled Key	4.3%
CWE-693	Protection Mechanism Failure	4.3%

Average number of vulns per project: 7.0

SOC 2 Pentest Findings – Attack Vector Distribution

Attack vector	% of SOC 2 findings	What this means
Authenticated user abusing permissions	~38%	Most serious issues require a login but no elevated privilege
Information leakage/reconnaissance	~25%	Systems expose data or behavior that helps attackers escalate
Session or token abuse	~14%	Weak session lifecycle enables takeover or persistence
Cryptographic / protection weakness	~13%	Controls exist but are incorrectly implemented
Unauthenticated external access	~7%	Fewer "open door" bugs than expected
Chained exploitation	~3%	Small number, but typically highest impact

Executive Analysis – SOC 2–Driven Assessments

The top CWEs identified during SOC 2–driven assessments are overwhelmingly centered on access control, data exposure, and enforcement of security controls, indicating that the systems tested are predominantly multi-user, application-centric platforms such as SaaS products and internal web services. The most frequent weaknesses (authorization bypass, improper access control, session handling flaws) do not reflect a lack of perimeter security, but rather failures inside trusted application boundaries.

Attack-vector analysis reinforces this picture. Nearly two-thirds of SOC 2 findings require some level of authentication, with the dominant risk being legitimate users performing actions or accessing data beyond their authorization. This highlights that identity is present, but authorization logic is incomplete, inconsistent, or incorrectly enforced, particularly at the object or tenant level.


A significant portion of findings also fall into information leakage and reconnaissance vectors, meaning systems unintentionally reveal internal details that reduce attacker effort and increase exploit reliability. While these issues are often perceived as low impact in isolation, they materially enable chaining and escalation when combined with access-control weaknesses.

Overall, the findings show that SOC 2–scoped environments generally have basic security controls in place, but struggle with the correct, consistent implementation of those controls across applications and services. The primary risk is not external compromise, but misuse of legitimate access leading to data exposure or privilege escalation, which directly undermines the trust guarantees SOC 2 is intended to provide.

ISO 27001 Penetration Tests

The assessments triggered by ISO 27001 compliance needs consisted of web application penetration testing and web & API security assessments – the surfaces ISO 27001 programs tend to validate when proving risk treatment and control effectiveness.

ISO 27001 Findings by Severity



Top 10 findings

CWEN	Name	% of ISO findings
CWE-284	Improper Access Control	12.9%
CWE-200	Exposure of Sensitive Information to an Unauthorized Actor	8.6%
CWE-209	Generation of Error Message Containing Sensitive Information	7.1%
CWE-327	Use of a Broken or Risky Cryptographic Algorithm	4.3%
CWE-693	Protection Mechanism Failure	4.3%
CWE-203	Observable Discrepancy	4.3%
CWE-345	Insufficient Verification of Data Authenticity	4.3%
CWE-613	Insufficient Session Expiration	4.3%
CWE-601	URL Redirection to Untrusted Site (Open Redirect)	4.3%
CWE-639	Authorization Bypass Through User-Controlled Key	4.3%

Average number of vulns per project: 11.67

ISO 27001 – Attack Vector Distribution

Attack vector	% of ISO findings	Interpretation
Authenticated user abusing permissions	~42%	Core risk lies within authenticated application usage
Information leakage/reconnaissance	~23%	Systems expose details that facilitate further attacks
Cryptographic / protection weakness	~15%	Controls exist, but are implemented inconsistently
Session or token abuse	~12%	Weak session lifecycle and token handling
Unauthenticated external access	~6%	Few perimeter-level exposures
Chained / multi-step exploitation	~2%	Rare, but highest potential impact

Executive Analysis – ISO 27001 Findings & Attack Vectors

The ISO 27001–driven assessments reveal a security posture dominated by application-level and identity-centric risks, rather than exposed perimeters or missing controls. The top CWEs (improper access control, authorization bypass, information disclosure, weak cryptographic practices) indicate that most organizations pursuing ISO 27001 have implemented security controls, but struggle with correct enforcement and consistency across systems.

Attack-vector analysis shows that the primary risk materializes after authentication. Nearly half of all ISO findings can be exploited by a legitimate user abusing permissions or roles, pointing to systemic weaknesses in authorization design, tenant separation, and business-logic enforcement. This aligns with ISO 27001's risk-based model, where failures often occur not from lack of controls, but from incomplete translation of risk treatment into application behavior.

Information leakage and cryptographic weaknesses form a secondary but important layer of risk. These issues rarely cause compromise on their own, but they reduce attacker effort, increase exploit reliability, and enable chaining, especially when combined with authorization flaws. The relatively low presence of unauthenticated attack vectors suggests that baseline perimeter security is generally in place, and that ISO-scoped environments are failing primarily within trusted boundaries.


In summary, ISO 27001 assessments expose the same fundamental challenge seen in SOC 2: security controls exist, but their enforcement breaks down in real application workflows. The dominant risk is not external attack, but misuse of legitimate access leading to data exposure or privilege escalation, undermining the assurance ISO 27001 is meant to provide.

PCI DSS Penetration Tests

PCI DSS-driven penetration tests represented a focused subset of the 2025 assessment portfolio. In contrast to broader web or infrastructure assessments, PCI testing scopes were typically narrower and control-oriented, concentrating on systems within or connected to the Cardholder Data Environment (CDE). As a result, the findings profile reflects a higher concentration of access control, authentication, and cryptographic weaknesses rather than large volumes of application logic flaws.

Although the overall number of findings per project was lower compared to general assessments, over one-third of PCI findings were classified as High or Critical, underscoring that even tightly scoped compliance-driven tests continue to surface materially impactful security weaknesses.

PCI DSS Findings by Severity



Top 10 findings

CWEN	Name	% of PCI findings
CWE-327	Use of a Broken or Risky Cryptographic Algorithm	14.3%
CWE-284	Improper Access Control	10.7%
CWE-200	Exposure of Sensitive Information to an Unauthorized Actor	7.1%
CWE-307	Improper Restriction of Excessive Authentication Attempts	7.1%
CWE-653	Improper Isolation or Compartmentalization	7.1%
CWE-311	Missing Encryption of Sensitive Data	7.1%
CWE-319	Cleartext Transmission of Sensitive Information	7.1%
CWE-522	Insufficiently Protected Credentials	7.1%
CWE-326	Inadequate Encryption Strength	7.1%
CWE-798	Use of Hard-coded Credentials	7.1%

Average number of vulns per project: 3.1

PCI DSS – Attack Vector Distribution

Attack vector	% of PCI findings	Interpretation
Cryptographic / protection weakness	~25%	Weak or outdated cryptography directly affects protection of cardholder data and is a core PCI DSS concern
Information leakage/reconnaissance	~18%	Systems expose sensitive details (headers, error messages, configurations) that aid further compromise
Session or token abuse	~18%	Session handling and token lifecycle weaknesses create risk of account or transaction hijacking
Authenticated user abusing permissions	~11%	Excessive privileges or missing role separation allow misuse within trusted contexts
Unauthenticated external access	~7%	Limited number of true perimeter failures, suggesting baseline exposure is mostly controlled
Other/chained exploitation paths	~21%	Findings that require combination of misconfigurations or do not map cleanly to a single vector

PCI Executive Analysis

The PCI-driven assessments in this dataset primarily consist of application security testing, infrastructure and network penetration testing, and PCI segmentation validation, reflecting the practical focus of PCI on isolating cardholder data environments and validating boundary controls. This mix indicates that PCI testing is not limited to application flaws, but also targets network exposure and trust boundaries, where segmentation failures can have immediate compliance impact.

The most common findings cluster around authentication and access control weaknesses, insecure configurations, and cryptographic failures, highlighting persistent challenges in protecting sensitive payment data at both the application and transport layers. The presence of cryptographic issues suggests that while encryption is often implemented, it is not always applied correctly

or consistently, particularly in key management, protocol selection, or certificate handling.

Attack vector analysis shows that most exploitable weaknesses are reachable via network-accessible and application-layer vectors, meaning attackers would not require privileged access to begin exploitation. Although Critical findings are relatively rare, the prevalence of High and Medium issues indicates realistic opportunities for attack chaining, where multiple smaller weaknesses could be combined to compromise cardholder data.

Overall, the data suggests that PCI programs are effective at preventing obvious systemic breakdowns, but operational gaps – especially around segmentation enforcement, identity controls, and cryptographic hygiene – remain the dominant sources of residual risk in real-world PCI environments.

9. Conclusion

This report brings together the results of hundreds of penetration testing engagements to provide a consolidated view of the security weaknesses observed across modern production environments.

The vulnerabilities discussed throughout the report were identified through manual penetration testing of real systems, under real constraints, and within realistic threat models. As such, the findings reflect not hypothetical risks, but concrete failure modes that repeatedly surface when systems are examined from an adversarial perspective.

Penetration testing plays a distinct role in the security lifecycle. Unlike automated scanning, it exposes how systems behave when controls are misused, combined, or bypassed in ways that were not anticipated during design. Many of the weaknesses highlighted in this report are difficult to detect automatically and are often missed by control-based assessments, yet they have a disproportionate impact on real-world security outcomes.

Across industries and technologies, the report shows that security risk is shaped less by the presence of individual controls and more by how consistently and correctly they are applied across environments. This makes penetration testing particularly valuable as a diagnostic tool that identifies where security assumptions break down in practice.

The aggregated results presented here are intended to support security teams in prioritizing effort, validating assumptions, and focusing on the areas where testing consistently reveals meaningful risk. Used as a reference point, this report can help organizations better understand how their security posture compares to broader trends, and where they can apply defensive effort most effectively.



Blaze Information Security is a leading provider of penetration testing and offensive security services, delivering comprehensive security assessments across web applications, APIs, mobile applications, cloud infrastructure, and specialized technologies. Through VulnKeep, our Penetration Testing as a Service (PTaaS) platform, we provide clients with real-time vulnerability visibility, audit-ready reporting, and integrated remediation workflows.

For more information about penetration testing services or to discuss your organization's security assessment needs, contact Blaze Information Security.

www.blazeinfosec.com