AMPCUS CYBER

INTELLIGENT CYBERSECURITY DELIVERED

BEYOND THE ANNUAL PEN TEST

Continuous AI-Powered Exploit Validation for Financial Services

AMPCUS CYBER
INTELLIGENT CYBERSECURITY DELIVERED

SECTION 01

EXECUTIVE SUMMARY

Financial institutions have spent the past decade building some of the most disciplined security testing programs of any industry. Annual penetration tests are scheduled, scoped, executed, and reported. Regulators are satisfied, audit committees sign off, and yet cyber incidents against the sector have continued to climb sharply: financial-sector cyber incidents more than doubled year over year, from 864 in 2024 to 1,858 in 2025, representing roughly one in five attacks recorded worldwide.[1]

The gap is not effort but timing. A penetration test is a photograph of an environment at a single moment. Financial services environments, by contrast, are in constant motion: new APIs are published to support fintech partnerships, mobile banking releases ship on weekly cycles, and core systems are continuously reconfigured. Between the date a penetration test concludes and the date the next one begins, the environment it once described no longer exists.

Meanwhile, the attacker's side of the equation has accelerated even faster. Research tracking the gap between vulnerability disclosure and first exploitation found that window collapsing from an average of 745 days in 2020 to just 44 days by 2025,[2] and AI-assisted exploit development is compressing it further still. A testing model anchored to a calendar date was never built to operate against a threat model measured in days.

This whitepaper examines why the point-in-time penetration test can no longer serve as the sole proof of resilience, what regulators including the EU's DORA and the PCI Security Standards Council now expect instead, and how continuous, AI-powered exploit validation, the model behind Mirror, closes the gap between when risk appears and when it is proven, not assumed.

1,858
Financial-sector cyber incidents recorded in 2025.

2.2×
Year-over-year increase from 864 incidents in 2024.

1 in 5
Share of all attacks worldwide now hitting finance.

> *The question is no longer whether your institution can pass next year's penetration test. It is whether it could survive an attack launched the day after this year's test ended.*

AMPCUS CYBER
INTELLIGENT CYBERSECURITY DELIVERED

IN THIS WHITEPAPER

WHAT WE COVER

01 Why the annual pen test model is structurally mismatched to financial services risk.

02 What DORA, PCI DSS 4.0.1, and SWIFT CSCF require and what they don't.

03 A real-world pattern: from an exposed API to a core banking compromise.

04 The continuous validation lifecycle behind ComplyX Mirror.

05 Board-level KPIs for measuring proof and posture.

AMPCUS CYBER
INTELLIGENT CYBERSECURITY DELIVERED

SECTION 02

THE WIDENING ATTACK SURFACE

From Digital Transformation to Risk Transformation

Banking has moved decisively beyond the branch and the core ledger. Open banking initiatives, embedded finance partnerships, mobile-first customer experiences, and API-driven fintech integrations have become the primary channels through which financial institutions compete. Each is, by design, a new connection point between systems that were never meant to be exposed and an ecosystem of partners, aggregators, and third-party developers now sitting on the other side of it.

The scale of that shift is difficult to overstate. The global open banking market, valued at roughly $39.8 billion in 2025, is projected to grow to $287.3 billion by 2033, a more than fourfold increase, as regulatory mandates and consumer demand for embedded financial services accelerate adoption worldwide.[3]

Every API endpoint added to support that growth is a new permutation of the attack surface. Banking and financial services recorded the highest year-over-year growth in API vulnerability attacks of any sector tracked, and API-related attacks against financial institutions have grown by more than 400%.[4] Akamai's 2025 API Security Report found 88.7% of financial firms experienced an API-related incident in the prior year the highest rate of any sector, at an average cost above $830,000 once downtime, legal exposure, and remediation are factored in.[5]

Digital transformation without continuous validation does not remove risk, it relocates and multiplies it. The institutions most exposed today are not the ones moving slowest on digital channels. They are the ones whose testing programs have not kept pace with how quickly those channels expand.

$287.3B

Projected open banking market by 2033, up from $39.8B in 2025.

400%+

Growth in API-related attacks against financial institutions.

88.7%

Of financial firms hit by an API-related incident in the past year.

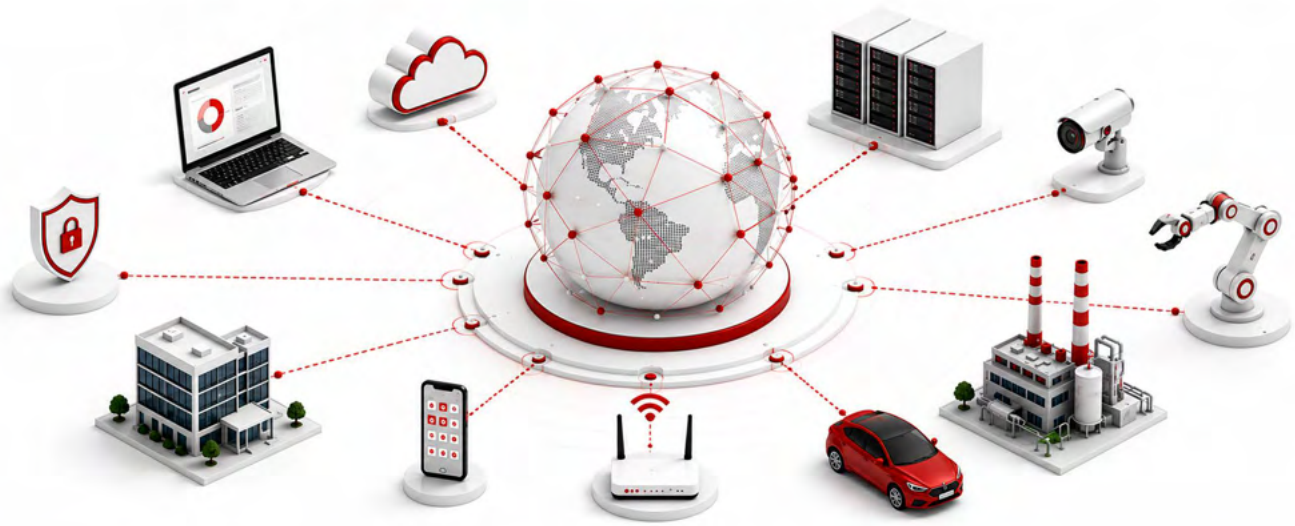

AMPCUS CYBER
INTELLIGENT CYBERSECURITY DELIVERED

SECTION 03

AI-ACCELERATED THREATS

Attackers have automated what defenders still schedule once a year

While financial institutions have largely kept testing a manual, calendar-driven discipline, attackers have spent the last two years automating theirs. Research analyzing tens of thousands of disclosed vulnerabilities found that AI-assisted exploit development compressed the average time between a vulnerability's disclosure and a working exploit from 125.3 days in January 2025 to under twelve hours by April 2026.[6] Separately, VulnCheck found that 32.1% of known exploited vulnerabilities in the first half of 2025 showed evidence of exploitation on or before the day the CVE was formally published up from 23.6% the year before.[7]

This is not a marginal improvement in attacker tooling. It is a structural change in what "fast enough" means for a defender. AI gives adversaries three capabilities a once-a-year testing model has no mechanism to counter:

- **Automated reconnaissance at machine speed.** AI-driven scanning fingerprints newly deployed APIs, subdomains, and cloud assets within hours of them going live long before they would ever appear in a pentest scoped against last year's asset inventory.
- **Automated exploit-chain construction.** Generative models help stitch together individually low-severity findings into a working multi-step attack path, the kind of chaining that traditional vulnerability scanning is structurally unable to demonstrate.
- **Continuous probing, not periodic scanning.** Where a human-led assessment runs for a defined window, automated attacker tooling runs continuously against the entire internet-facing footprint of an institution, every day of the year.

125 days to <12 hrs
AI-compressed time from disclosure to working exploit, Jan 2025 vs. Apr 2026

32.1%
Of known exploited vulnerabilities hit on or before public disclosure

Fraud teams see a parallel acceleration: deepfake-enabled fraud losses against financial institutions exceeded $410 million globally in the first half of 2025 alone, with generative-AI-enabled fraud projected to approach $40 billion annually by 2027.[8] The common thread is the same, AI is already a force multiplier for adversaries. The best-positioned institutions use it as a force multiplier for validation, not just detection after the fact.

AMPCUS CYBER
INTELLIGENT CYBERSECURITY DELIVERED

SECTION 04

WHY THE ANNUAL MODEL FALLS SHORT

The traditional model rests on an assumption that no longer holds: that the system tested in week one will still resemble the system in production in week fifty. A scoped, manual engagement produces a report, a remediation plan, and in most institutions, a long pause before the next test begins. It was built for environments that changed slowly. It is now asked to certify environments that change daily.

Testing Benchmark	Cadence	Detail
PCI DSS Penetration Test	**Every 12 months**	Minimum required cadence for internal / external pentests under Requirement 11.4.
DORA Threat-Led Test (TLPT)	**Every 36 months**	Minimum cadence for advanced TLPT at systemically significant institutions.
Digital Release Cadence	**Daily**	Benchmark cadence for elite-performing software delivery teams across digital channels.
Time-to-Exploit	**~44 days (2025)**	Average window between disclosure and active exploitation of a vulnerability.

Set side by side, the mismatch is structural rather than incidental. A control validated as effective twelve months ago says little about whether it still holds after dozens of intervening releases. The limitations compound in predictable ways:

- **Point-in-time blind spot.** A clean report dated last quarter says nothing about a feature shipped last week.
- **Manual throughput can't match velocity.** A multi-week engagement can't repeat often enough for weekly, let alone daily change.
- **Scope drift outpaces scope definition.** New APIs and acquired systems go live outside the prior engagement's boundary.
- **Findings without proof.** Scanning stops at "may be exploitable," leaving prioritization to severity scores, not demonstrated impact.

None of this argues against the formal, regulator-recognized penetration test, it remains a compliance necessity. It argues that the formal test can no longer be the **only** *mechanism an institution relies on to know where it stands.*

SECTION 05

EVOLVING COMPLIANCE EXPECTATIONS

DORA, PCI DSS 4.0.1, and Beyond

Regulators have not missed the cadence mismatch they are actively legislating against it. The direction across every major framework governing financial services testing points the same way: toward proof performed on live systems, not documentation prepared in advance of an audit.

The Digital Operational Resilience Act (DORA)

In full application across the EU financial sector since January 2025, DORA requires in-scope entities to perform standard ICT security testing at least annually, while systemically significant institutions must additionally undergo Threat-Led Penetration Testing (TLPT) at least once every three years under Article 26. TLPT must be performed on live production systems, must include relevant third-party ICT providers in scope, and must use external threat intelligence with one in every three cycles using fully external testers. The objective is not to audit the entire IT estate, but to demonstrate the ability to withstand a realistic, intelligence-led attack.

PCI DSS v4.0.1

The current operative standard requires internal and external penetration testing of the cardholder data environment at least annually and after any significant change, under Requirement 11.4. Institutions relying on segmentation to reduce scope must validate it annually; service providers, every six months. Quarterly vulnerability scans are required in addition to the annual test. All future-dated 4.0 requirements became mandatory in March 2025, every assessment today is evaluated against the full requirement set.

SWIFT Customer Security Controls Framework (CSCF)

SWIFT's CSCF requires participating institutions to undergo periodic independent assessment of the controls protecting their SWIFT-connected environment, with the framework's scope and control set updated on a recurring basis to reflect emerging attack patterns.

Taken together, these frameworks describe a clear intent: institutions should be able to demonstrate, under realistic conditions, that their controls work. What none solve on their own is the gap **between** *testing cycles precisely where unvalidated risk accumulates.*

SECTION 06

FROM EXPOSED API TO CORE COMPROMISE

The following composite reflects a pattern commonly observed across financial services environments, not a specific named incident and shows how the gap described earlier plays out in practice.

A regional bank prepares to launch an account-aggregation partnership with a fintech. To meet a committed date, engineering stands up a new API endpoint exposing balance and transaction-history data. It inherits a broader identity and access role than the integration requires a shortcut taken under deadline pressure, to be tightened after launch. The bank's most recent annual pen test, completed months earlier, never saw this endpoint; it did not exist when the engagement was scoped.

Within days, automated reconnaissance identifies the newly registered subdomain. The flaw is unremarkable: Broken Object Level Authorization (BOLA), consistently the most common vulnerability class in reported API incidents. By manipulating an account identifier, the attacker retrieves data for accounts well beyond the one the partner was authorized to query.

From there the path widens. Session tokens harvested through the endpoint are valid against the bank's broader digital banking API. Lateral movement continues toward payment-initiation workflows the multi-step chain a scanner reports as disconnected findings, but an attacker experiences as one continuous path.

> *From a leadership perspective, this is not an API misconfiguration story. It is a timing story. Every control that ultimately failed had a documented owner and a clean prior penetration test report.*

The lesson is not that the program was negligent. It is that its primary proof point an annual, scope-bound penetration test was structurally incapable of seeing a risk introduced after the test concluded.

SECTION 07

THE EXPLOIT VALIDATION MATURITY CURVE

Most institutions have functioning visibility and assessment programs. Far fewer have closed the gap between identifying a weakness and proving what it would cost the business if exploited. Maturity in exploit validation tends to progress across four recognizable stages:

STAGE 01
Visibility
What assets, APIs, and partner connections exist across the environment.

STAGE 02
Assessment
Identifying vulnerabilities and misconfigurations through scanning and review.

STAGE 03
Validation
Proving which findings are exploitable in context, and how far they chain.

STAGE 04
Continuous Assurance
Validation as an ongoing program tied to every change, evidence retained.

INCREASING MATURITY

Most institutions plateau between stage two and stage three: they generate long, severity-ranked lists of findings without ever confirming which an attacker could realistically use, in what order, and against what business impact. That plateau is comfortable, it produces a report but it leaves the central question regulators, boards, and auditors increasingly ask unanswered: not *"what did you find,"* but **"what did you prove."**

Institutions operating at stage four treat exploit validation the way they treat fraud monitoring or transaction surveillance, a continuous operational discipline rather than a periodic compliance exercise. That shift, more than any single tool, is what separates organizations that can demonstrate resilience on demand from those that can only demonstrate it once a year.

SECTION 08

QUANTIFYING THE COST OF UNVALIDATED RISK

The financial impact of operating on assumptions rather than proof is no longer theoretical. It shows up directly in breach-cost data, regulatory exposure, and ransom economics across the sector:

$5.56M
Average cost of a financial-sector data breach in 2025, second only to healthcare.

241
Days, on average, from intrusion to full breach containment, 2025.

745 to 44
Days: collapse in the disclosure-to-exploitation window, 2020 vs. 2025.

1,858
Cyber incidents against the financial sector in 2025, double 2024's 864.

Every additional day a vulnerability sits unvalidated is a day an institution cannot answer a simple question with confidence: is this exploitable in our environment, today? That uncertainty has a price. Ransomware operators targeting financial services demanded a record-high median ransom of $3 million in the latest reporting period, with exploited vulnerabilities cited as the most common technical root cause of successful attacks.[14] Regulatory exposure compounds the cost: DORA, NYDFS 23 NYCRR 500, GLBA, and SEC disclosure rules each layer notification, audit, and remediation obligations on top of direct incident response.

The deeper inefficiency is where remediation budget goes when findings are not validated. Teams that prioritize purely by CVSS severity routinely spend weeks remediating theoretical risks while exploitable, business-critical paths go unaddressed. Focusing remediation on confirmed exploitability rather than theoretical severity directs a constrained budget toward the small subset of findings that matter, lowering the effective total cost of a continuous validation program well below the cost of either an unvalidated breach or a remediation program built on guesswork.

AMPCUS CYBER
INTELLIGENT CYBERSECURITY DELIVERED

SECTION 09

CRITICAL USE CASES FOR CONTINUOUS VALIDATION

Effective exploit validation in financial services must address risk across every layer where digital banking, payments, and partner ecosystems intersect.

Digital & Mobile Banking Testing

Continuous validation of customer-facing web and mobile apps as releases ship, so features are proven secure before, not months after, they reach customers.

API & Open Banking Validation

Testing for BOLA, excessive data exposure, and weak authentication across partner-facing and internal APIs as they are added.

Payment & PCI Scope Validation

Exploit validation across cardholder data environment components and segmentation boundaries, supporting PCI DSS 11.4 with continuous evidence.

Core Banking & SWIFT Infrastructure

Validating exposure across systems connected to payment messaging infrastructure, supporting SWIFT CSCF assessment expectations.

Cloud & Hybrid Infrastructure

Validating misconfigurations, excessive permissions, and identity exposures across the multi-cloud and hybrid environments underpinning banking platforms.

Third-Party & Fintech Chaining

Testing how a compromise originating in a single partner integration could chain laterally into core systems.

Pre-Release & CI/CD-Integrated Validation

Embedding exploit validation into the release pipeline itself, so new code and new API endpoints are tested for exploitability before or immediately after reaching production, rather than waiting for the next scheduled engagement.

SECTION 09

SECURING OPEN BANKING & THE API ATTACK SURFACE

Open banking deserves particular attention because it concentrates several of the risks described elsewhere in this whitepaper into a single layer. Banking and financial services recorded the highest year-over-year growth in API vulnerability attacks of any sector tracked, and Salt Security's research found that nearly one-third of financial services APIs carried at least one critical vulnerability.[15]

~1 in 3

Financial services APIs carry at least one critical vulnerability.

28.5%

Of institutions report a complete inventory of their APIs and which return sensitive data.

Compounding the problem, a majority of institutions cannot fully answer the most basic question of what their open banking attack surface even includes, which endpoints exist, and which expose sensitive data.

Continuous, API-aware exploit validation directly addresses this layer: as new endpoints are published to support partner integrations, they are discovered, tested for authorization and data-exposure flaws, and validated for exploitability on an ongoing basis, closing the visibility gap that traditional, calendar-bound testing was never built to track.

SECTION 10

THE CONTINUOUS VALIDATION LIFECYCLE

Operationalizing exploit validation effectively requires a repeatable cycle, not a one-time project. ComplyX Mirror is built around four continuous phases:

01
Discover
Continuously map the external and internal attack surface new APIs, endpoints, and partner integrations as they appear.

02
Chain
Autonomous AI agents simulate real attacker behavior, linking individual weaknesses into multi-step attack paths.

03
Validate
Each chain is tested against real environments to confirm exploitability and capture proof-of-exploit evidence, not a score.

04
Prioritize & Remediate
Findings ranked by demonstrated business impact a remediation order grounded in evidence, not severity alone.

↻ **A closed loop, not a line.** Prioritize & Remediate continuously feeds back into Discover newly added APIs and releases are picked up automatically, and configuration changes are revalidated as they happen.

This is the structural shift that distinguishes continuous exploit validation from a faster version of the same annual test. An institution's security posture evolves alongside its release cadence rather than its audit calendar. It is not a bigger snapshot, it is a fundamentally different operating model.

SECTION 11

THE COMPLYX MIRROR ADVANTAGE

ComplyX Mirror is Ampcus Cyber's AI-powered penetration testing platform, purpose-built to close the validation gap. Rather than generating another list of theoretical findings, Mirror deploys autonomous AI agents that discover, chain, and validate vulnerabilities across web applications, APIs, infrastructure, and mobile apps simulating real attacker behavior to deliver proof-of-exploit evidence in place of guesswork. Mirror Intelligence extends that visibility outward, giving a continuous view of external and third-party exposure.

- **Continuous, not calendar-bound.** Validation aligned to release cycles, not a once-a-year engagement window.
- **Proof over probability.** Every finding demonstrated as exploitable in context, not flagged by severity score alone.
- **Full-stack coverage.** Web, mobile, API, and infrastructure tested as one interconnected attack surface.
- **Built for regulated environments.** Evidence structured to support PCI DSS 11.4, DORA obligations, and audit requests.

ONE CONNECTED RESILIENCE PROGRAM

ComplyX GRACE	ComplyX Mirror	ComplyX Wizard
Centralizes continuous compliance evidence.	Proves exploitability with evidence.	Extends visibility into third-party risk.

GRACE
From Chaos to Continuous Compliance

mirror
Reflect. Detect. Protect.

Wizard
Intelligence beyond the perimeter

Stop Guessing. Start Proving.

A connected view across exploitability, compliance, and ecosystem risk not three disconnected reports.

SECTION 12

BEYOND THE SCAN: 2026 BOARD-LEVEL KPIS FOR CYBER RESILIENCE

Reporting findings counts to the board tells leadership very little about actual resilience. The following KPIs reframe testing outcomes around proof and business impact, the language boards and regulators increasingly expect.

KPI	Measurement Focus	Strategic Outcome
Validated Control Efficacy	% of internet- and partner-facing assets with confirmed, non-exploitable findings	Assurance that controls hold under real attack conditions
Mean Time to Validate (MTTV)	Time from new API/release to first exploit-validation pass	Closes the window between deployment and assurance
Exploitable Attack-Path Density	Confirmed multi-step exploit chains across web, API, and partner systems	Visibility into systemic risk beyond isolated CVEs
Remediation Velocity (Proof-Based)	Time to remediate findings by demonstrated exploitability	Directs limited capacity to risks with real business impact
Audit & Regulatory Readiness	% of PCI DSS 11.4 & DORA obligations supported by continuous evidence	Reduces audit scramble; strengthens defensibility

These KPIs enable a shift from technical reporting to business-aligned security metrics that hold up in front of a board, examiner, or regulator asking for evidence rather than assurances.

SECTION 13

THE PATH FORWARD

From Periodic Assurance to Continuous Proof. The threat landscape, regulatory direction, and pace of digital release in financial services have all moved past what an annual penetration test was ever designed to certify. Attack surfaces will keep expanding as open banking grows, exploitation timelines will keep compressing as attacker tooling becomes automated, and regulators will keep raising the evidentiary bar for what counts as proof of resilience.

The question financial security leaders now face is no longer whether vulnerabilities exist in their environment. It is whether they can prove, **continuously**, which of those vulnerabilities an attacker could use, and demonstrate that proof on demand not just once a year.

For CISOs and Risk Leaders: Three Priorities

- Move from periodic testing to continuous, evidence-based exploit validation across web, mobile, API, and infrastructure layers.
- Integrate validation into the release pipeline itself, so new code and new APIs are tested before they accumulate into the next year's blind spot.
- Treat proof-of-exploit evidence as a board- and regulator-facing asset, not a security team artifact that never leaves the SOC.

Stop Guessing. Start Proving.

Ampcus Cyber helps financial institutions move from periodic assurance to continuous, evidence-based exploit validation with ComplyX Mirror.

AMPCUS CYBER
INTELLIGENT CYBERSECURITY DELIVERED

SOURCES

REFERENCES

1. Check Point Research, 2025 Finance Threat Landscape Report — financial-sector cyber incidents rose from 864 (2024) to 1,858 (2025): https://blog.checkpoint.com/research/the-three-most-disruptive-cyber-trends-impacting-the-financial-industry-today/

2. Flashpoint, N-Day Vulnerability Trends: The Shrinking Window of Exposure (2025/2026) — average time-to-exploit fell from 745 days (2020) to 44 days (2025): https://flashpoint.io/blog/n-day-vulnerability-trends-turn-key-exploitation/

3. Grand View Research, Open Banking Systems Market Size & Share Report — global market valued at $31.61B in 2024, projected to reach $135.17B by 2030 at a 27.6% CAGR: https://www.grandviewresearch.com/industry-analysis/open-banking-systems-market

4. IBM Threat Intelligence research on API-related attack growth in financial services, referenced via: https://www.instantpay.in/blog/api-banking-security-best-practices/

5. Akamai, 2025 API Security Report — 88.7% of financial firms experienced an API-related incident; 28.5% report a full API inventory, referenced via: https://www.openbankingexpo.com/usa/open-finance-open-apis-and-the-rising-stakes-of-security/

6. Cogent Research, The Detection Gap: How Exploits Are Outpacing Scanners (2026) — AI-compressed time-to-exploit from 125.3 days (Jan 2025) to under 12 hours (Apr 2026): https://www.prnewswire.com/news-releases/cogent-research-exploits-outpace-scanner-detection-for-62-of-critical-vulnerabilities-as-ai-compresses-time-to-exploit-to-under-12-hours-302783104.html

7. VulnCheck, State of Exploitation: 1H-2025 — 32.1% of KEVs exploited on/before CVE disclosure date: https://www.vulncheck.com/blog/state-of-exploitation-1h-2025

8. Fourthline, Deepfakes in Financial Services: How AI Fraud Is Reshaping Risks in 2026 — H1 2025 deepfake fraud losses and 2027 projection: https://www.fourthline.com/blog/deepfakes-in-financial-services

9. PCI Security Standards Council, PCI DSS v4.0.1, Requirement 11.4 (penetration testing and segmentation testing cadence): https://blog.pcisecuritystandards.org/just-published-pci-dss-v4-0-1

10. European Union, Digital Operational Resilience Act (DORA), Article 26 — Threat-Led Penetration Testing requirements: https://www.digital-operational-resilience-act.com/Article_26.html

11. Google Cloud / DORA Research Program, Accelerate State of DevOps findings on elite-performer deployment cadence, referenced via: https://www.scrums.com/blog/deployment-frequency-benchmarks

12. OWASP, API Security Top 10 (2023), API1:2023 Broken Object Level Authorization (BOLA) ranked the most critical API security risk: https://owasp.org/API-Security/editions/2023/en/0xa1-broken-object-level-authorization/

13. IBM Security, Cost of a Data Breach Report 2025 — financial sector breach cost ($5.56M) and global breach lifecycle (241 days): https://www.ibm.com/reports/data-breach

14. Invenio IT, Ransomware in Financial Services: 2026 Insights — record median ransom demand and root-cause data: https://invenioit.com/continuity/ransomware-attacks-finance/

15. Salt Security research on API vulnerability prevalence in financial services, referenced via: https://www.instantpay.in/blog/api-banking-security-best-practices/

AMPCUS CYBER

INTELLIGENT CYBERSECURITY DELIVERED

About Ampcus Cyber

Ampcus Cyber is a leading global cyber security organization headquartered in Chantilly, Virginia. We are dedicated to providing comprehensive, cutting-edge solutions to protect your digital assets. Founded with a mission to combat the ever-evolving cyber threats, we combine expertise, technology, and a client-centric approach to deliver unmatched security services.

Ready to take the next step?

Connect with our cybersecurity experts today.

+1 (703) 310-6237

letsconnect@ampcuscyber.com

www.ampcuscyber.com