

# The Seam Report 2026

Where intrusions bend between the systems you test. Cyber, physical and human attack paths in the United States and the Gulf.

Edition  
**October 2026**

Data as of  
**September 28, 2026**





# The Seam Report 2026

Where intrusions bend between the systems you test. Cyber, physical and human attack paths in the United States and the Gulf.

|             |                                                                                                                   |
|-------------|-------------------------------------------------------------------------------------------------------------------|
| Edition     | October 2026                                                                                                      |
| Data as of  | September 28, 2026                                                                                                |
| Publisher   | Bendings, <a href="https://bendings.com">bendings.com</a>                                                         |
| How to cite | Bendings. The Seam Report 2026. October 2026. <a href="https://bendings.com/research">bendings.com/research</a> . |

This report contains no client data. Every number in it is computed by Bendings from public data or quoted from a named primary source, and each source is listed in the notes.

## About Bendings

Bendings is a cybersecurity firm founded in 2013. It runs reconnaissance, red team, penetration testing, physical and social-engineering testing, blue team support and incident response for governments, banks and asset managers in the United States and the Gulf Cooperation Council states. Its own AI model, Bendings Gauge, runs continuous red and blue team work inside limits each client approves, on Bendings Frame, the firm's own operating system for digital and physical security operations.

# Contents

|          |                                                                                                                                                                                                                                                                            |           |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
|          | <b>Why a firm that works by referral wrote this</b>                                                                                                                                                                                                                        | <b>5</b>  |
|          | A note on what this report is, what it is not, what we sell and how to check it.                                                                                                                                                                                           |           |
|          | <b>Key findings</b>                                                                                                                                                                                                                                                        | <b>6</b>  |
|          | Nine findings from public evidence. Each gives its window and denominator, and the chapter named after it shows the work.                                                                                                                                                  |           |
| <b>1</b> | <b>Intrusions bend at the seams between tests</b>                                                                                                                                                                                                                          | <b>8</b>  |
|          | A method for writing down attack paths so that the places they cross from one kind of system, one owner or one authority to another become visible and countable.                                                                                                          |           |
| <b>2</b> | <b>Tests cover the steps and miss the joins</b>                                                                                                                                                                                                                            | <b>12</b> |
|          | Seventy-eight public intrusions from 2013 to 2026, written step by step in bend notation. Almost every step falls where some test looks, but fewer than half of the paths could be followed from first step to last.                                                       |           |
| <b>3</b> | <b>The controls at the edge are now a way in</b>                                                                                                                                                                                                                           | <b>18</b> |
|          | Among recently published flaws, products sold to keep intruders out rose from 11 to 18 percent of those CISA confirmed as exploited, comparing the first nine months of 2022 and 2026. They reach the list within days and carry the ransomware mark more often than most. |           |
| <b>4</b> | <b>Whoever resets a password is part of identity</b>                                                                                                                                                                                                                       | <b>27</b> |
|          | Help desks, call centers, support contractors and hiring teams issue and restore access every day. Adversaries have learned to ask them for it, pay them for it or apply for it.                                                                                           |           |
| <b>5</b> | <b>Supplier seams are crossed often, tested rarely</b>                                                                                                                                                                                                                     | <b>33</b> |
|          | Counting each supplier's breach once, about one in five incidents US companies disclosed came through a supplier, as did one in three intrusions in our corpus. Of 41 US and Gulf testing regimes, 9 put supplier-run systems in scope.                                    |           |
| <b>6</b> | <b>Few tests cross the door, camera or valve</b>                                                                                                                                                                                                                           | <b>39</b> |
|          | The network ends at a door, a camera, a cash machine and a valve. Network monitoring and most testing rules stop short of each of these seams, and adversaries have crossed all four.                                                                                      |           |
| <b>7</b> | <b>In the Gulf, a cloud region has a street address</b>                                                                                                                                                                                                                    | <b>46</b> |
|          | In 2026 drone strikes took down cloud facilities in the UAE and Bahrain, and some data held only there was lost for good. Gulf rules that keep data and testers in the country now sit against recovery plans that assume another region.                                  |           |
| <b>8</b> | <b>AI agents under test reached real organizations</b>                                                                                                                                                                                                                     | <b>54</b> |
|          | Adversaries use AI mostly to go faster, and defenders use it to find flaws first. In 2026 AI systems under cyber evaluation crossed from the test into real organizations, and no testing regime we read sets rules for AI that tests.                                     |           |

|    |                                                                                                                                                                                                                                                             |     |
|----|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 9  | <b>Testing rules stop at the seams attackers cross</b>                                                                                                                                                                                                      | 62  |
|    | Of 41 US and Gulf testing regimes, one requires a physical intrusion test, five a test of people and nine a test that reaches systems a supplier runs. None of the 50 regimes we read, including nine foreign benchmarks, sets rules for AI-driven testing. |     |
| 10 | <b>Gulf deadlines fall due hours before US ones</b>                                                                                                                                                                                                         | 70  |
|    | Jurisdiction is a seam, and reporting deadlines are where it bites. Three Gulf central banks want a first report within one hour of a serious incident. US bank regulators allow 36 hours from the bank's own determination.                                |     |
| 11 | <b>Banks guard mail but rarely say where to report</b>                                                                                                                                                                                                      | 79  |
|    | We measured two seams that anyone can see from outside at 159 US and Gulf regulators and banks. The email seam is mostly closed. The reporting seam, where an outsider with a flaw to report looks for the right person, is mostly unmarked.                |     |
| 12 | <b>Test the crossings, not only the systems</b>                                                                                                                                                                                                             | 85  |
|    | Keep a register of seams, measure how many have ever been crossed in a test, design scenarios around bends, and arrange the authority to cross them before anyone needs it.                                                                                 |     |
|    | <b>How Bendings works at the seams</b>                                                                                                                                                                                                                      | 89  |
|    | The one place in this report where we describe ourselves and what we sell. Everything here is a standing commitment, not a claim about results.                                                                                                             |     |
| A  | <b>Method and limitations</b>                                                                                                                                                                                                                               | 90  |
|    | How each dataset was built, what it can and cannot show, and where to find the data and code.                                                                                                                                                               |     |
| B  | <b>Bend notation, on one page</b>                                                                                                                                                                                                                           | 93  |
|    | A reference card for writing any intrusion, test scenario or incident timeline as a path of domains and owners.                                                                                                                                             |     |
| C  | <b>A seam register to start from</b>                                                                                                                                                                                                                        | 95  |
|    | Eighteen crossings from the public record, each with a documented example, ready to be given owners and a test date.                                                                                                                                        |     |
| D  | <b>The incident corpus</b>                                                                                                                                                                                                                                  | 98  |
|    | All 78 cases, in order of year, with each path in bend notation and the sources it was coded from.                                                                                                                                                          |     |
| E  | <b>Glossary</b>                                                                                                                                                                                                                                             | 101 |
|    | Terms this report defines, and the outside terms it relies on, in the sense we use them.                                                                                                                                                                    |     |
|    | <b>Notes</b>                                                                                                                                                                                                                                                | 104 |

# Why a firm that works by referral wrote this

A note on what this report is, what it is not, what we sell and how to check it.

---

Bendings has worked by referral since 2013. Most clients find us through other clients, and the work stays where it is done. This report contains nothing from a client engagement. It has no counts, no anecdotes dressed up as statistics and no figures that begin with "in our experience".

We wrote it anyway, because the pattern we are hired to find does not show up in any single organization's data. A service desk records a password reset. The identity team records a new authenticator. The cloud team records an administrator signing in from somewhere new. Each record is ordinary. Put in order, they are an intrusion. This report is about the seams between those records.

We set out to test one claim, that intrusions cross from one kind of system to another, and from one owner to another, in ways that the scope of common tests and the rules for testing do not follow. We used public evidence only. That meant 78 intrusions from 2013 to September 2026, coded step by step from government advisories, court records, regulators' decisions and the victims' own filings, with a random sample of the better-documented cases coded again by an analyst who had not seen the first coding. It meant the full US catalog of vulnerabilities known to be exploited, 135 cybersecurity incidents US public companies disclosed on Form 8-K from December 18, 2023 to September 28, 2026, 50 sets of testing rules, 96 incident notification duties, and a passive measurement of 159 banks and authorities in the United States and the Gulf.

Tests cover the steps of documented intrusions and miss the joins between them, most often where a step runs through another organization that the victim cannot authorize a tester to touch. Where the evidence did not support us, each chapter says so in a section of its own. The email seam we expected to find open was mostly closed. Our

corpus is too small to show whether intrusions through other organizations are becoming more common. Artificial intelligence appears in about one percent of the business email compromise losses reported to the FBI for 2025.

Weigh one fact about us as you read. Bendings sells the testing this report finds missing, including red team, physical, social-engineering and AI-driven testing, and some of our proposals would bring firms like ours more work if adopted. The proposals run from a notation for attack paths and a register of seams to controls for AI systems that test and changes to testing rules. Each is labeled as ours and kept apart from the evidence.

The notes list each source with its address and the date we read it. The case files, derived data and code are published with the report; raw measurements that would name individual banks or ransomware victims are held back.

The Gulf appears in these pages more than in most reports of this kind, because we work there and because 2026 changed it. In March, drones struck two cloud facilities in the United Arab Emirates, and a strike nearby damaged one in Bahrain. In September the provider said it could not restore data that customers had held only in its Bahrain region or only in one of its UAE zones. Recovery plans that assumed a cloud region could not be lost as a unit now have a dated counterexample.

Bendings  
October 2026

# Key findings

Nine findings from public evidence. Each gives its window and denominator, and the chapter named after it shows the work.

2 days

median from publication to CISA's exploited list for 80 flaws in products sold as security controls, published 2022 to 2025. For all 651 flaws, 10 days.

Chapter 3

29%

of new Form 8-K cyber incidents were first filed as material after the SEC staff statement of May 21, 2024 (33 of 112), down from 74% before it (17 of 23).

Chapter 5

1 hour

to first report a serious incident to the central bank in Bahrain, Qatar and Kuwait. US bank regulators allow 36 hours.

Chapters 7 and 10

**01 Tests cover the steps. They miss the joins.**  
In 70 of 78 public intrusion paths from 2013 to 2026, every step on the victim's own systems fell in a domain some single-domain test covers. A typical red team scope could follow 33 paths end to end. Of the 45 it could not, 38 ran through another organization the victim cannot authorize a tester to touch. Chapter 2.

**02 Products sold to keep intruders out are a growing way in.**  
Among recently published flaws, security controls rose from 10.7 percent (13 of 122) of CISA's exploited-list additions in January to September 2022 to 17.6 percent (36 of 204) in the same months of 2026. The 37 added so far in 2026 exceed any full year since 2022. Chapter 3.

**03 Most new SEC incident filings no longer call the incident material.**  
After May 2024, 33 of 112 new incidents were first filed under Item 1.05, against 17 of 23 before. Chapter 5.

**04 Whoever resets a password is part of identity.**  
US advisories from 2023 to 2026 keep naming help-desk staff who reset credentials when a caller asks. Coinbase paid \$311.2 million in cash in 2025 on an incident that began with bribed support staff. Chapter 4.

**05 Harm reaches further than the adversary goes.**  
In 7 of 78 cases, operational technology stopped because IT was lost, against 6 in which the adversary acted in it. In 48, harm reached other owners. Chapters 2 and 6.

**06 In the Gulf, recovery advice and residency rules pull apart.**  
After the March 2026 drone strikes, AWS told customers to recover "ideally in Europe", and later said some data held only in its Bahrain region or one UAE zone could not be restored. Qatari and new UAE rules keep core records in-country. Chapter 7.

**07 AI agents under test reached real organizations.**  
In July and August 2026, OpenAI, Anthropic and the UK AI Security Institute each reported AI systems under cyber evaluation reaching organizations no one had pointed them at. Chapter 8.

**08 The fix, not the finding, takes the time.**  
The PHP project fixed two critical flaws on October 21, 2026. It had seen the first signs of one in 2024, and work on a fix stalled for months. The next day, 32 UK government services were unavailable and four partly so, largely as a precaution, with no confirmed compromise of central government systems. Chapter 8.

## 09 Testing rules stop at the building, the person and the supplier.

Of 41 US and Gulf testing regimes, 1 requires a physical intrusion test, 5 a test of people and 9 a test of supplier-run systems. None of 50 sets rules for AI-driven testing. Chapter 9.

## What cut against us

- Bend counts follow from our coding rules. They measure how many control failures the sources document, not whether intrusions bend.
- The corpus is too small to show a trend in owner seams. They appear in 9 of 21 paths from 2013 to 2018, 10 of 18 to 2022 and 19 of 39 since.
- The email seam is mostly closed. Of 159 banks and authorities we measured, 147 enforce DMARC.
- Most disclosed US incidents, 99 of 135, state no third-party path, and the 36 that do are at most 28 supplier events.
- Corrected for CISA's 2022 backlog, the security-control rise is about two thirds, not the near tripling in the raw figures.

- The Gulf share of ransomware leak-site posts rises under a broad count, but much of that rise comes from the tracker and the strict count shows no steady trend, so it is not a finding here.
- Faster Gulf clocks are rules, not results. We found no data showing Gulf institutions detect intrusions faster.

## What we propose

These are our proposals, not findings. Bend notation for writing any intrusion or test as a path of domains and owners (chapter 1, Appendix B). A seam register with a named owner on each side of every crossing (chapter 12, Appendix C). Five board measures of seams covered, owned, detected and tested (chapter 12). Seam-led test scenarios, with the authority to cross arranged in advance (chapter 12). Ten controls for AI systems that test (chapter 8). And testing rules aimed at the bodies that can change them (chapter 9).



# 1

## Intrusions bend at the seams between tests

A method for writing down attack paths so that the places they cross from one kind of system, one owner or one authority to another become visible and countable.

|                                                          |    |
|----------------------------------------------------------|----|
| Four terms, used the same way throughout                 | 10 |
| Why seams, and not domains                               | 10 |
| Bend notation records the domain and owner of every step | 10 |
| A bend diagram gives each domain its own lane            | 11 |
| A bend count measures documented control failures        | 11 |

On February 21, 2025, the people who approve transfers from the Bybit exchange's Ethereum cold wallet signed what their screens showed as a routine movement of funds. The transaction they signed was not the one they saw. It handed control of the wallet to North Korean operators, who took about \$1.5 billion.<sup>1,2</sup>

Bybit's own systems were not where the theft began. Weeks earlier, a developer at Safe{Wallet}, the supplier whose web interface Bybit used to approve transactions, ran a malicious project on a work computer. The attackers used that developer's live cloud session to enter the supplier's cloud account, and there they changed the code the interface served so that it would tamper only with transac-


tions from Bybit's wallet.<sup>3,2</sup> Then they waited for Bybit's staff to do their jobs.

Written as a path, the theft has four steps. A person at the supplier, the supplier's cloud identity, the supplier's cloud storage, the victim's wallet. Three steps belonged to one company and the last to another. No test that either company could have commissioned on its own systems would have followed it from start to finish.

Case 1 Bybit cold wallet theft via Safe{Wallet}

2025 · Finance · Global · TraderTraitor (North Korea; also tracked as Lazarus Group, UNC4899), state · Theft

H<sub>s</sub> / I<sub>s</sub> / C<sub>s</sub> | / V



- 1 H Got a Safe{Wallet} developer to run a malicious project on their work Mac, likely through social engineering<sup>3,2</sup> supplier

2 I Used the developer's live cloud session token to enter Safe{Wallet}'s AWS account, sidestepping multi-factor authentication after failing to register a new MFA device<sup>3,2</sup> supplier
- 3 C Altered the web code Safe{Wallet} served from its cloud storage so it would tamper only with transactions from Bybit's wallet<sup>3,2,4</sup> supplier

4 V Bybit's signers approved a routine cold-to-warm wallet transfer that the tampered interface displayed as legitimate but that handed control of the wallet to the attackers, who then moved the funds<sup>1,2,4</sup>

About \$1.5 billion in virtual assets (over 400,000 ETH and stETH) was taken from a Bybit Ethereum cold wallet on February 21, 2025, per the FBI and Bybit. Bybit's own systems showed no sign of compromise; the malicious transaction came from its wallet supplier's web interface.

Confidence: medium. Discovered by the victim. Dwell time not stated in sources.

We call each change between steps a bend, and the boundary a bend crosses a seam.

## Four terms, used the same way throughout

### Seam

A boundary where the owner of a system, the controls that protect it, the team that watches it or the legal authority over it changes.

**Domain.** The kind of system or thing an adversary acts on in one step. We use nine: human, identity, edge, internal IT, cloud and SaaS, operational technology, physical, value transfer and AI systems. The test for which domain a step belongs to is practical. Ask which set of controls had to fail for the step to work.

**Owner.** Whose system, account or person it was: the victim, a supplier, a customer, an employee's personal account, or another party such as a telecom carrier or a domain registrar.

**Bend.** A change of domain or of owner between two consecutive steps of an intrusion.

**Seam.** The boundary a bend crosses.

We use these terms to count. With a domain and an owner on every step, bends can be tallied across dozens of cases, and questions usually argued from anecdote can be answered with a number and a denominator.

## Why seams, and not domains

Security organizations are built by domain, and for good reasons. Identity has its own tools, staff and maturity model. So do the network team, the facilities team and the supplier-risk function. Testing follows the same lines. A web application test has a list of addresses. A phishing exercise

has a list of recipients. A physical assessment has a list of doors. Each list is scoped, priced and signed on its own.

Adversaries work from a list of objectives and take whatever route reaches one. In 38 of the 78 documented intrusions in chapter 2, the route ran through another organization's systems or people.

Four things change at a seam, and each one helps the adversary.

- 1 **The control owner.** On each side of a seam someone is responsible for their half. The crossing itself is often nobody's.
- 2 **The logs.** Telemetry stops at ownership boundaries. A bank cannot read its supplier's sign-in records, and a help desk ticket does not sit next to the identity provider's log unless someone joined them on purpose.
- 3 **The authority.** A tester needs permission from every owner whose systems a scenario touches. A regulator's testing scheme stops at the licensed entity. A national law stops at the border.
- 4 **The assumption.** Each side tends to believe the other is watching the crossing.

## Bend notation records the domain and owner of every step

To count bends we needed to write intrusions down the same way every time. Bend notation records each step as a one-letter domain code, with a small letter when the step belonged to someone other than the victim:  $H_s$  is a person at a supplier,  $I_p$  an employee's personal account. A slash marks a change of domain and a vertical bar a change of owner.

| Code | Domain                 | The adversary is acting on                                                                                                       |
|------|------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| H    | Human                  | A person's judgment: a call to a help desk, a bribed support agent, a fake job applicant                                         |
| I    | Identity               | Accounts, credentials, tokens, keys and the processes that issue and reset them                                                  |
| E    | Edge                   | Systems exposed to the internet on purpose: VPN and firewall appliances, file transfer servers, public applications, DNS records |
| N    | Internal IT            | Workstations, servers, directory, hypervisors, backups, build systems                                                            |
| C    | Cloud and SaaS         | Cloud consoles and storage, SaaS tenants and the integrations between them                                                       |
| O    | Operational technology | Industrial control and safety systems, building systems, controllers                                                             |
| P    | Physical               | Premises, planted devices, removable media, radio proximity, hardware in transit                                                 |
| V    | Value transfer         | Payment and card switches, SWIFT terminals, trading and settlement, wallet signing                                               |
| A    | AI systems             | Models and agents with access to data, tools or code                                                                             |

The Bybit theft reads  $H_s / I_s / C_s | / V$ . The Marks & Spencer intrusion of April 2025 reads  $H_s | / I / N$ , and chapter 4 tells that story.<sup>5,6</sup> Two businesses, the same shape. In both, the first step was at a supplier, where the victim's own tests could not go.

Value transfer is its own domain because, at a bank or an exchange, the systems that move money are the objective and are often run by a separate team. AI systems are one because models with access to mail, code and tools are now places an adversary can act.

MITRE ATT&CK, the public knowledge base of adversary techniques, names what the adversary did at each step.<sup>7</sup> Bend notation adds which domain and whose systems. Our case files list ATT&CK identifiers where the mapping is clear, in 77 of the 78.

# A bend diagram gives each domain its own lane

Each case is drawn as in Case 1 above. The nine domains are horizontal lanes. A path runs left to right in the order the steps happened, level within a domain, climbing or falling where it bends. A blue line broken by a short upright rule has crossed into someone else's systems. A dashed start means the sources do not say how the adversary got in. The amber dot is the objective.

## Tests add, paths multiply

One test per domain is nine tests. The ordered pairs of domains an adversary can move between number 72, and each can cross an owner seam as well. Testing grows with the number of domains. The crossings grow with its square.

# A bend count measures documented control failures

We applied the notation to 78 publicly documented intrusions from 2013 to September 2026, and chapter 2 reports what the paths show. Counting their bends does not prove that intrusions bend. Under our coding rules a case needs at least two steps, a new step starts only when a different control had to fail, and consecutive steps in the same domain with the same owner are merged. Almost every step boundary is therefore a bend, and a path's bend count measures how many separate control failures the sources document. It rises with documentation. The mean is 2.48 bends in the 31 cases coded with high confidence, 2.18 in the 39 at medium and 1.0 in the 8 at low. Appendix B gives the rules and Appendix A their effects.

What the corpus can show is where the bends fall and whose systems they cross, and chapter 2 rests on those measures. Our claim is that intrusions cross kinds of system and owners in ways common test scopes and testing rules do not follow. Each chapter says where the record does not support it.



## 2

# Tests cover the steps and miss the joins

Seventy-eight public intrusions from 2013 to 2026, written step by step in bend notation. Almost every step falls where some test looks, but fewer than half of the paths could be followed from first step to last.

|                                                             |    |
|-------------------------------------------------------------|----|
| The steps are covered, the joins are not                    | 13 |
| Short paths across several domains                          | 14 |
| Identity is where paths turn                                | 15 |
| Nearly half the paths pass through someone else             | 16 |
| Consequences bend too                                       | 16 |
| Most intrusions were found by someone other than the victim | 17 |

We read the public record of 78 intrusions, from Target in 2013 to the attacks on Polish energy sites and the Stryker wiper in 2026, and wrote each one down in bend notation. Every step cites a public source: a government advisory, an indictment, a regulator's decision, the victim's own filing, the responder's report or, where nothing better exists, the press.

Thirty-five cases are in the United States and fifteen in the Gulf. Thirteen are in finance, fifteen in government and eight in energy. State services ran 32, criminal groups 33, and 12 are not attributed. Appendix D lists every path.

This is not a random sample. It is a sample of intrusions someone explained in public, which favors large victims, US disclosure rules and the cases responders chose to write up. We chose the cases, and we held the thesis while choosing them. The counts describe these 78 cases and nothing wider. As a scale for noise, the 95 percent Wilson interval around the share that passed through another owner, 38 of 78, runs from 38 to 60 percent. Slices such as eras or regions are smaller and noisier, so we treat differences between them as description.

### Checking the coding

A second analyst coded 16 cases, drawn at random from the 70 coded with high or medium confidence, without seeing the first coding. The two agreed in 16 of 16 on whether a path passed through another owner and whether it left a typical red team scope. They agreed on the exact sequence of domains in 8 of 16, and on whether a path touched identity in 13 of 16, the weakest agreement of any domain. Low-confidence cases were not rechecked. Appendix A gives the details.

## The steps are covered, the joins are not

The steps are covered. In 70 of the 78 paths, every known step on the victim's own systems falls in a domain that some single-domain test covers: an external or internal penetration test, a cloud configuration review, a phishing exercise, a physical test or an OT assessment. In 69, every such step falls inside a typical red team scope, which chains people, identity, the edge, internal systems, cloud and payment systems. Appendix A defines each test.


## 33 of 78

documented intrusion paths could be followed from first step to last by a typical red team scope, although in 69 every step on the victim's own systems fell inside it.

Bendings incident corpus, 2013 to September 2026

The joins are not. A typical red team scope could have followed 33 of the 78 paths from start to finish. The other 45 left its scope, 38 of them because a step ran through another organization's systems or people, which the victim cannot authorize a tester to touch. Six reached operational technology, four needed a physical step and one ran through an AI system. A tester can go only where every owner on the path has agreed. These paths kept going.

Figure 1 45 of 78 paths left a typical threat-led red team scope, most often through another organization




A purposive sample of well-documented cases, not a random sample of all intrusions. Counts describe these cases only. A path can have more than one reason. Typical scope: human, identity, edge, internal IT, cloud and value-transfer systems owned by the victim.

Source: Bendings incident corpus, 78 publicly documented intrusions, 2013 to September 2026, coded in bend notation from primary sources. Bendings analysis. n = 78

Single-domain tests are not designed to follow whole paths. Under our definitions they could have followed at most 9 of the 78, a count that runs from 4 to 28 under other reasonable definitions. It falls to 4 if undisclosed first steps count against the test, and rises to 28 if a phishing exercise

is credited with the payload it runs on a workstation and an external test with its first internal host. Following a path is also a different question from finding a weakness in it. Closing any one step breaks the path.

Figure 2 A typical red team scope could follow 33 of 78 documented paths from first step to last; single-domain tests at most 9



A purposive sample of well-documented cases, not a random sample of all intrusions. Counts describe these cases only. A path counts only if every known step falls in the test's domains and every step belongs to the victim. Unknown first steps are not held against a test. Counting them against it, the single-domain figure is 4; crediting phishing and external tests with an internal foothold, it is 28.

Source: Bendings incident corpus, 78 publicly documented intrusions, 2013 to September 2026, coded in bend notation from primary sources. Bendings analysis. n = 78


Leaving out the eight cases coded with low confidence changes little. Of the other 70, 36 pass through another owner and 43 leave a typical red team scope. The 13 finance cases look the same. Six pass through another owner, and a typical red team scope could have followed six. Every finance case from 2023 on is a crypto exchange or a payment processor. The most recent bank victim in the corpus is Capital One, in 2019.

## Short paths across several domains

After merging consecutive steps in the same domain and with the same owner, the median path has three steps. Forty of the 78 touch three or more domains, and seven stay in one. How many bends a path shows follows from our coding rules and from how much the sources say, as

chapter 1 explains. The findings here rest on where the bends fall and whose systems they cross.

Figure 3 Seventy-eight documented intrusions, each drawn as one path across the nine domains



Every path starts at the left edge and runs in the order its steps happened. Darker bundles are routes many cases share.

Source: Bendings incident corpus, coded from public primary sources, snapshot September 28, 2026. n = 78 cases

## Identity is where paths turn

Of the 151 changes of domain in the corpus, 98 run into or out of identity. The most common single bend is identity to internal IT, 30 times, where a stolen or forged credential is used to move inside. Next come internal IT back to identity (22), the edge to internal IT (15) and identity to cloud (15).

Figure 4 Two in three bends run into or out of identity

|                          | Human | Identity | Edge | Internal IT | Cloud and SaaS | Operational technology | Physical | Value transfer | AI systems |
|--------------------------|-------|----------|------|-------------|----------------|------------------------|----------|----------------|------------|
| ? Not disclosed          | 2     | 7        | 1    | 8           | 1              |                        |          |                |            |
| H Human                  | 1     | 7        |      | 11          |                |                        | 2        |                |            |
| I Identity               | 3     | 3        | 3    | 30          | 15             | 4                      |          | 2              |            |
| E Edge                   | 1     | 8        | 1    | 15          | 1              |                        |          |                |            |
| N Internal IT            | 4     | 22       | 1    | 12          | 2              | 1                      | 1        | 7              |            |
| C Cloud and SaaS         |       | 3        |      | 1           | 3              |                        |          | 1              |            |
| O Operational technology |       |          |      |             |                | 3                      |          |                |            |
| P Physical               |       | 1        |      | 2           |                |                        |          | 1              |            |
| V Value transfer         |       |          |      | 1           |                |                        |          | 2              |            |
| A AI systems             |       |          | 1    |             |                |                        |          |                |            |

Rows are the domain of a step, columns the domain of the step after it.

Rows are the domain of a step, columns the domain of the step after it. The diagonal counts steps where only the owner changed.

Source: Bendings incident corpus, coded from public primary sources, snapshot September 28, 2026. n = 195 step transitions

Internal IT also touches 98 of the 151, but mostly as the place an intrusion ends. It is the last step in 36 paths. Identity appears in 53 paths, is the last step in 3 and the first in 12, so it sits mostly in the middle. An organization that tests identity by reviewing its identity provider's configuration tests the turn without the steps on either side.

The share was similar in each era. It was 9 of 21 paths from 2013 to 2018, 10 of 18 from 2019 to 2022 and 19 of 39 from 2023 to 2026. One case moves an era's share by three to six points, so the corpus cannot show a trend in either direction. It does show that owner seams were common in documented cases from the start.

## Nearly half the paths pass through someone else

Thirty-eight of the 78 paths include a step on systems or people the victim did not own: a supplier, a customer, an employee's personal account, a carrier, a neighbor's Wi-Fi. In 24 the first documented step was outside the victim's estate. In nine more, the sources show the intrusion began at a supplier or customer without saying how, and we record that first step as undisclosed but owned by them. For Home Depot, Norsk Hydro and the Aramco contractor case, that inferred step is the only reason the path counts here. The most common change of owner is supplier to victim, 21 times. Target's 2013 path reads  $H_s | / I / N / V / N$ . A refrigeration contractor's network credentials were the way in to the network that ran the stores' payment terminals.<sup>8,9</sup> Chapter 5 takes up the supplier seam.

## Consequences bend too

In 9 of the 78 cases, harm reached a domain the adversary never touched, because the victim or someone downstream shut systems down or could not run without them. Seven of the nine reached operational technology, which chapter 6 examines.

- Colonial Pipeline halted all 5,500 miles of its pipeline in May 2021 as a precaution. Its chief executive's testimony describes no adversary action in operational systems.<sup>10</sup>
- Norsk Hydro's 35,000 employees worked by hand for about three weeks in 2019, at a cost of more than \$71 million.<sup>11</sup>
- Airports in Europe went back to manual check-in, gates and baggage when ransomware hit Collins Aerospace's

passenger processing systems in September 2025.<sup>12,13</sup>

In 48 of the 78 cases the harm reached other owners as well, most often customers. A seam register that records only how an adversary could cross is half a register. The other half records what stops working on the far side of each seam when the near side is lost.

## Most intrusions were found by someone other than the victim

Where the sources say who found the intrusion, 37 of 68 were found by someone other than the victim: the adversary announcing itself (14), a third party such as a researcher or journalist (12), a supplier (7) or law enforcement (4). For campaigns coded as one path, the discoverer is the campaign's, not one victim's. Only 14 cases give both the date of first access and the date of discovery. Their median is about a month, four are two days or less, and the figure is not comparable with the dwell times responders publish from their caseloads. The longest is 1,502 days at Starwood, where the intrusion began in 2014, two years before Marriott bought the company, and was flagged by a contractor in 2018.<sup>14</sup>

## The Gulf cases rest on thinner sources

The fifteen Gulf cases rest on different kinds of source than the US cases, so the two compare only loosely. Seven Gulf paths begin at the edge, against 13 of the 63 elsewhere, and five pass through another organization, among them the card processors behind the 2012 and 2013 cash-outs against RAKBANK and Bank Muscat.<sup>15,16</sup> With 15 cases, one case moves a share by almost seven points. Six Gulf cases are coded with low confidence, four from press accounts alone, because Gulf authorities publish few incident narratives. Chapter 7 returns to the Gulf.

## Where the data disagrees

Nine paths could have been followed end to end by one single-domain test, and four of them are fully documented. At Colonial Pipeline and Change Healthcare, a stolen creden-

tial opened a remote access service without multi-factor authentication, and the intruders moved inside.<sup>10,17</sup>

Microsoft's 2024 intrusion was a password spray against a legacy test tenant account, followed by access to corporate mail.<sup>18</sup> The Iranian changes to controller logic at US sites in 2026 never left operational technology.<sup>19</sup> The other five qualify only because the sources do not say how the adversary got in. Some intrusions are simple, and a good penetration test would have found the door.

Part of identity's weight comes from our notation. Appendix B codes a harvested password, and a stolen password used on a VPN, as steps in identity. A notation that folded credential use into the next domain would show fewer identity bends, and identity is the domain on which our two coders agreed least.

None of this says single-domain testing is useless. Most documented paths needed more than one kind of test, and a permission the victim could not give alone.

## What to do

- **Heads of offensive security:** for each red team scenario, write the intended path in bend notation before the test, and record in the report every seam the team reached but was not authorized to cross.
- **Heads of identity:** treat the processes that issue, reset and delegate credentials as part of the identity perimeter, and join help desk, identity provider and cloud sign-in records so a single path can be followed across them.
- **Supplier risk leads:** for the ten suppliers whose people or credentials can reach your systems, find out whether anyone has ever tested the path from their side to yours.
- **Operations and resilience leads:** list what stops beyond each IT dependency, including what you would halt yourselves, and rehearse that decision.



# 3

## The controls at the edge are now a way in

Among recently published flaws, products sold to keep intruders out rose from 11 to 18 percent of those CISA confirmed as exploited, comparing the first nine months of 2022 and 2026. They reach the list within days and carry the ransomware mark more often than most.

|                                                          |    |
|----------------------------------------------------------|----|
| Security products are a growing share of exploited flaws | 20 |
| Security products reach the catalog within days          | 21 |
| Ransomware follows the controls                          | 22 |
| Why the edge is a seam                                   | 23 |
| Prediction scores rarely flag exploited flaws in advance | 24 |
| CISA now assumes the box may be compromised              | 25 |

On January 10, 2024, CISA added two flaws in Ivanti Connect Secure, a remote access gateway, to its catalog of vulnerabilities known to be exploited. The National Vulnerability Database (NVD) published them two days later.<sup>20,21</sup> At MITRE, the adversary had been inside since December 31, 2023.<sup>22</sup>


By MITRE's account, the two flaws let the intruders plant a web shell on the gateway and get past multi-factor authentication. They took over sessions and credentials, reached the NERVE research network, used administrator credentials on its VMware servers, built virtual machines of their

own and took data out on January 19, 2024. They kept reading files until the middle of March.<sup>22,23</sup> In bend notation it is E / I / N. The first step was on a product bought to keep intruders out.

Case 2 MITRE NERVE research network

2024 · Defense · US · Overlaps with UNC5221, state · Espionage

E / I / N



- Victim's systems

Someone else's systems

Owner seam

Domain bend, 3-4-5 slope

Objective
- 1 E

Exploited two Ivanti Connect Secure zero-days (CVE-2023-46805, CVE-2024-21887) and planted the ROOTROT web shell, bypassing MFA<sup>22</sup>
- 2 I

Hijacked sessions and credentials to log in to NERVE accounts over RDP<sup>22</sup>
- 3 N

Used compromised administrator credentials on vCenter and ESXi, created rogue VMs, deployed BRICKSTORM and BEEFLUSH, and exfiltrated data on 2024-01-19 with file access continuing to mid-March<sup>22,23</sup>

The adversary gained access to MITRE's NERVE research and prototyping network, maintained persistence in its VMware infrastructure and exfiltrated data.

Confidence: high. Discovered by the victim. Dwell time not stated in sources.

This chapter asks how common that first step has become. It rests on our analysis of the CISA Known Exploited Vulnerabilities catalog, version 2026.09.27, which lists 1,728 flaws. We count the 1,417 added from January 1, 2022 to September 27, 2026, leaving out the 287 loaded on its first day, November 3, 2021.<sup>24</sup> We map each product to a category and two flags, whether it faces the internet by design and whether its vendor sells it as a security control.

The mapping and its 89 debatable calls are published with the data.

## Security products are a growing share of exploited flaws

Products sold as security controls include firewalls, VPN gateways, email and web security gateways, endpoint protection, network access control, privileged access management and their consoles. From January 1 to September 28, 2026, CISA added 37 flaws in them, already more than in any full year from 2022 to 2025 (34, 13, 24 and 35). They were 15.2 percent of 244 additions. In the same months of 2022 the share was 5.4 percent, 28 of 521.

Part of that jump is an artifact. In 2022 CISA was still loading old flaws, and 72.4 percent of its January to September additions had been published two or more years earlier, against 14.3 percent in 2026. Counting only flaws published in the year of listing or the year before, security controls went from 10.7 percent (13 of 122) to 17.6 percent (36

of 204), rising every year in between. That is a rise of about two thirds, not a near tripling.


# 17.6%

of recently published flaws added to the KEV catalog from January 1 to September 28, 2026 were in products sold as security controls, 36 of 204. In the same months of 2022 the share was 10.7%, 13 of 122.

CISA KEV catalog, NVD, Bendings analysis <sup>20,21</sup>

Of the 143 security-control additions since 2022, 99 are network edge appliances. Products built to face the internet, a wider group that adds mail servers, managed file transfer, remote support servers and the like, moved the same way. Counting recent flaws only, their share went from 29.5 percent (36 of 122) in 2022 to 38.7 percent (79 of 204) in 2026. Moving every debatable product to the other side of the line leaves the 2026 share of all additions between 32.0 and 40.6 percent.

**Figure 5** Counting only recently published flaws, products sold as security controls took 18% of KEV additions in January to September 2026, up from 11% in the same months of 2022



Restricted to flaws published in the year of listing or the year before. With CISA's 2022 backlog left in, the security-control share runs from 5.4 to 15.2 percent.


Restricted to KEV additions whose CVE was published in the year of addition or the year before, which removes CISA's backlog loads of old CVEs. Security controls: 13 of 122 in 2022, 36 of 204 in 2026. Denominators: 2022 n=122, 2023 n=100, 2024 n=97, 2025 n=142, 2026 n=204. The edge and security-control lines share their 2022 and 2026 values by coincidence of counts. The two groups overlap but are not the same products.

Source: CISA Known Exploited Vulnerabilities catalog, version 2026.09.27 (released 2026-09-27), retrieved 2026-09-28. Bendings analysis; product categories are Bendings' own mapping. CVE publication dates from the NVD CVE API 2.0. n = 665

# Security products reach the catalog within days

For the 651 flaws published from 2022 to 2025 and listed by September 27, 2026, the median time from NVD publication to listing was 10 days.<sup>21</sup> Network edge appliances took a median 5 days (80 flaws). Products sold as security controls took 2 days (80 flaws), and 37.5 percent were listed on or before the day they were published.

Figure 6 Edge appliance flaws reached the KEV a median 5 days after publication, against 10 days overall



CVEs published 2022-01-01 to 2025-12-31 and added to the KEV by 2026-09-27. Categories with at least 10 such CVEs. Per-category n: Desktop and server operating systems and drivers 111, Office, document and file-handling software 16, Mobile OS, chipsets and mobile apps 77, Browsers and browser plug-ins 36, Network edge and infrastructure appliances 80, Managed file transfer and file sharing 24, Security products (endpoint, email and web security, WAF, security management) 22, Backup and storage 16, Email and collaboration servers 54, Virtualization and hypervisors 11, Web application platforms, CMS and business web apps 51, Remote access and IT management tools 57, Developer tools, libraries and supply-chain components 28, Consumer and SOHO network devices and IoT 22, Data, middleware, analytics and AI platforms 19, Physical security devices (cameras, recorders, access control) 11. Negative lags occur when CISA added a CVE before NVD published it. A CVE added later would lengthen these medians.

Source: CISA Known Exploited Vulnerabilities catalog, version 2026.09.27 (released 2026-09-27), retrieved 2026-09-28. Bendings analysis; product categories are Bendings' own mapping. CVE publication dates from the NVD CVE API 2.0. n = 651

A short lag mostly means the flaw was disclosed because it was already being exploited, so it reflects how zero-days

are disclosed as much as how fast attackers move. Operating systems and browsers, at 0 and 2 days, are as


fast. The difference, in our reading, comes after listing. A browser updates itself. A firewall update waits for a network team, a restart of the device that carries all the traffic and often a change window. We did not measure that exposure.

---

## Ransomware follows the controls

CISA marks each entry where it knows of use in a ransomware campaign. Across all 1,417 additions since 2022 the mark is on 19.5 percent (277). On products sold as security controls it is on 35.7 percent, 51 of 143. Only backup and storage products, at 51.4 percent (18 of 35), carry it more often.

**Figure 7 Products sold as security controls were tied to ransomware in 36% of their KEV additions since 2022, against 20% for all additions**



CISA can add the ransomware mark after listing, so older entries have had longer to collect it. "Unknown" means CISA does not know of ransomware use, not that there was none.

CISA's knownRansomwareCampaignUse field, as of the snapshot, for additions 2022-01-01 to 2026-09-27. CISA can set the flag after addition, so older entries have had longer to be flagged. 'Sold as a security control' is a flag across categories (firewalls, VPNs, secure gateways, endpoint security and similar). Categories with at least 15 additions. n: All KEV additions 1417, Internet-facing by design 405, Sold as a security control 143, Backup and storage 35, Managed file transfer and file sharing 28, Email and collaboration servers 95, Desktop and server operating systems and drivers 249, Remote access and IT management tools 93, Network edge and infrastructure appliances 187, Developer tools, libraries and supply-chain components 74, Security products (endpoint, email and web security, WAF, security management) 39, Virtualization and hypervisors 15, Web application platforms, CMS and business web apps 109, Data, middleware, analytics and AI platforms 52, Browsers and browser plug-ins 146, Office, document and file-handling software 68, Consumer and SOHO network devices and IoT 58, Physical security devices (cameras, recorders, access control) 21, OT, ICS and building systems 18, Mobile OS, chipsets and mobile apps 114.

Source: CISA Known Exploited Vulnerabilities catalog, version 2026.09.27 (released 2026-09-27), retrieved 2026-09-28. Bendings analysis; product categories are Bendings' own mapping. n = 1,417

The corpus shows how. Boeing Distribution's 2023 intrusion is E / I / N. LockBit 3.0 affiliates read valid session cookies out of a NetScaler appliance's memory and re-played one.<sup>25</sup> With a working session, the crew needed no password and no second factor, so the sign-in checks never ran.

## Why the edge is a seam

An edge appliance is a seam, as chapter 1 defines it, in a single box, for three reasons.

It is a control and an entry point at once. Its owners test its rules and penetration testers test its exposure. Few programs ask what an adversary can do from inside it, because from inside it is still, on paper, the defense.

It holds keys to identity. A VPN gateway authenticates users, so it sees passwords, session tokens and often a directory service account. In our corpus of 78 documented intrusions (chapter 2), the edge was the first known step in 20 paths and led next to internal IT 15 times and to identity 8 times. At MITRE, at Boeing Distribution and in the Salt Typhoon campaign against telecom carriers, the next step was identity. For Salt Typhoon the joint advisory describes the campaign's pattern, not one carrier's path. The actors went after TACACS+ authentication servers and used service and automation credentials to reach other network devices.<sup>26</sup>


Its evidence often stays on the box. Appliances run closed operating systems, and many keep detailed logs locally, rotate them quickly and cannot run endpoint agents. We read this as why the appliance's owner and detection's owner are often different teams, with a gap an intruder can sit in. That is a judgment. The public record does not show how often appliance logs were missing when responders needed them.

---

## Prediction scores rarely flag exploited flaws in advance

The Exploit Prediction Scoring System, run by FIRST, estimates the chance that a published flaw will be exploited in the next 30 days, and many organizations patch by it.<sup>27</sup> We pulled each KEV addition's score for the day before CISA listed it, for the 843 additions from March 7, 2023 to September 27, 2026.<sup>28</sup>

Figure 8 The day before CISA listed them, 73% of scored KEV additions had an EPSS score below 0.1



EPSS changed model versions during the window, so scores across years are not strictly comparable.

KEV additions from 2023-03-07 to 2026-09-27: 843 additions. 611 had a score for the day before addition, or up to 3 days earlier when that daily file was missing (7 cases). 232 had no score yet; 225 of those were published on or after the target day. Median EPSS 0.0038; internet-facing median 0.0034. EPSS model versions in the window: v2023.03.01 (290), v2025.03.14 (243), v2026.06.15 (75), v2022.01.01 (3). Categories with at least 15 scored additions. Per-category scored n: Mobile OS, chipsets and mobile apps 42, Network edge and infrastructure appliances 69, Desktop and server operating systems and drivers 61, Email and collaboration servers 56, Managed file transfer and file sharing 20, Web application platforms, CMS and business web apps 69, Office, document and file-handling software 16, Physical security devices (cameras, recorders, access control) 19, Browsers and browser plug-ins 31, Consumer and SOHO network devices and IoT 30, Remote access and IT management tools 61, Security products (endpoint, email and web security, WAF, security management) 22, Data, middleware, analytics and AI platforms 34, Developer tools, libraries and supply-chain components 42.

Source: CISA Known Exploited Vulnerabilities catalog, version 2026.09.27 (released 2026-09-27), retrieved 2026-09-28. Bendings analysis; product categories are Bendings' own mapping. EPSS scores from the FIRST EPSS API ([api.first.org](https://api.first.org)), historical scores for the day before each addition. n = 611.

Of all 843, 611 had a score that day, and 447 of those (73.2 percent) were below 0.1. Network edge appliances sit at the high end of that range, not outside it. Of 117 edge additions, 61 of the 69 scored (88.4 percent) were below 0.1, with a median of 0.0023. The other 48 (41.0 percent) had no score, 47 of them because they had been published that day or later. Of 109 security-control additions, 51 of 59 scored were below 0.1 and 50 were unscored. Backup and storage (92.3 percent) and mobile (all 42 scored) sat lower still.

This is not a verdict on EPSS. A low score the day before listing does not prove the model wrong about the next 30 days. It shows that ranking a patch queue by prediction alone would not have put most exploited flaws near the top

when it mattered. At the edge, exploitation often arrives with the disclosure, and a model needs a published flaw to score.

## CISA now assumes the box may be compromised

Binding Operational Directive 26-04, dated June 10, 2026, changes what federal civilian agencies must do with entries CISA marks for forensic triage.<sup>29</sup>



**The text "forensic triage" means that the agency must complete remediation or mitigation action within the timeline (three days) and carry out a forensic triage of the asset to assess whether the system is compromised.**

CISA, Binding Operational Directive 26-04<sup>29</sup>

The first entry with the flag was added on July 1, 2026. From then to September 27, CISA asked agencies to check for compromise, not only to patch, on 69 of 98 additions (70.4 percent), including 32 of 39 internet-facing additions (82.1 percent) and 37 of 59 others (62.7 percent). Three months is a short record.

## Where the data disagrees

Internet-facing products as a group are not faster to the catalog. Their median was 11.5 days (206 flaws), against 9 days for the rest (445 flaws), slowed by web platforms, mail servers, remote access tools and home routers. The speed belongs to security controls and edge appliances, not to exposure as such.

The network edge category's own share of all January to September additions barely moved, from 12.9 percent in 2022 (67 of 521) to 15.6 percent in 2026 (38 of 244). Five vendors hold 65.8 percent of the 187 network edge additions since 2022, and Cisco's count is inflated by 2022 backlog loads.

Severity scores do not rescue a simple rule. Since 2022, 162 of 1,417 additions (11.4 percent) have an NVD base score below 7.0, as do 13.3 percent of internet-facing ones. Patching only flaws scored 7 and above would miss about one in nine known to be exploited.

The catalog itself has limits. It lists only flaws with reliable evidence of exploitation, a CVE identifier and a clear fix.<sup>30</sup> It records what CISA confirmed in a year, not what attackers did. And not every edge path starts with a flaw. Capital One in 2019 began with a misconfigured web application firewall, which no vulnerability catalog would list.<sup>31</sup>

## What to do

- **CISO:** keep a register of every product that both enforces security and accepts connections from the internet, with the credentials each holds, and review it as a list of entry points, not of controls.
- **Head of vulnerability management:** patch KEV-listed flaws in edge and security products within three days of listing, the clock BOD 26-04 sets federal agencies for entries flagged for forensic triage, whatever their CVSS or EPSS score, and start the check for compromise the same day.
- **Head of detection and response:** ship appliance logs to storage the appliance cannot alter, keep them for the weeks before a flaw becomes public, and join them with identity provider sign-ins so a session that starts at the gateway can be followed.
- **Head of identity:** give each appliance the narrowest directory account it needs, and after any exploitation rotate every credential it held and revoke the sessions it issued.
- **Head of offensive security:** add an assumed-compromise scenario that starts on the VPN or firewall and records how far the team gets toward identity and internal IT.
- **Procurement:** require vendors of edge and security products to support log export, integrity checking and a documented forensic triage procedure.



# 4

## Whoever resets a password is part of identity

Help desks, call centers, support contractors and hiring teams issue and restore access every day. Adversaries have learned to ask them for it, pay them for it or apply for it.

|                                                      |    |
|------------------------------------------------------|----|
| Every warning names the same desk                    | 28 |
| Access has a market price                            | 29 |
| Hiring became a way to issue credentials             | 29 |
| Most losses reported to the FBI ran through a person | 31 |
| Where the data disagrees                             | 32 |
| What to do                                           | 32 |

On September 14, 2023, Caesars Entertainment told the Securities and Exchange Commission that suspicious activity in its network came from "a social engineering attack on an outsourced IT support vendor used by the Company." The actor left with a copy of the loyalty program database.<sup>32</sup> Nothing in the filing describes a software flaw. Someone at a supplier's help desk did what the caller asked.

The people who reset passwords, enroll authenticator devices, look up customer records and approve new hires are the working parts of an identity system. They sit in the service desk, the contact center, an outsourced support firm or human resources. The identity team usually owns the directory and the sign-in policy. It rarely owns the phone line, and it almost never owns the hiring process.

In our corpus of 78 documented intrusions, a human step appears in 23 paths and is the first step in 16. The bend from human to internal IT occurs 11 times and the bend from human to identity seven times. The split between those two counts says more about disclosure than about the intrusions. Caesars' filing names the help desk and the network, and says nothing about the account in between, so we coded what the source says.

## Every warning names the same desk

The joint advisory on Scattered Spider, first issued by CISA, the FBI and partners on November 16, 2023 and updated on July 29, 2025, describes a group that "targets large companies and their contracted information technology (IT) help desks."<sup>33</sup>



**The threat actors conduct [spearphishing] calls to convince IT help desk personnel to reset passwords and/or transfer MFA tokens.**

CISA and FBI advisory AA23-320A<sup>33</sup>

The warnings since then repeat the point. On June 27, 2025 the FBI told the airline sector that the same group convinces help desk services "to add unauthorized MFA devices to compromised accounts" and targets "large corpo-

rations and their third-party IT providers."<sup>34</sup> On September 12, 2025 an FBI FLASH described callers posing as IT support who tricked customer support employees into connecting an attacker's tool to their firm's Salesforce instance.<sup>35</sup> In June 2024 the FBI and HHS warned that callers held stolen personal data, passed help desk identity checks and then diverted payments.<sup>36</sup>

The New York State Department of Financial Services sent the CISOs of its regulated entities an industry letter on February 6, 2026. It describes the same trick run in the other direction, with callers posing as the help desk to staff, and it asks firms to review who can change MFA enrollment.<sup>37</sup> Twitter in 2020 is the corpus case for that direction, H / I / N / H / N | / I<sub>c</sub> | / H<sub>t</sub>. Callers posing as the IT help desk collected VPN credentials and MFA approvals, then used employees' access to internal account tools to take over 130 user accounts.<sup>38</sup>

# 11%

of intrusions Mandiant investigated in 2025 began with voice phishing, second only to exploits at 32%. Email phishing fell to 6%. Mandiant's summary does not state the engagement count.

M-Trends 2026<sup>39</sup>

Mandiant's figure is one vendor's caseload, and it matches the advisories.<sup>39</sup> By August 2026 Google was tracking another vishing group, UNC6671, working against financial and legal firms, including private equity firms and rating agencies. Its pretext was an "urgent helpdesk mandate to enable FIDO2 passkeys."<sup>40</sup> The rollout of the strongest authenticator became the cover story for stealing the employee's sign-in.

The failure point in all of these documents is the same. A person with authority to change someone's credentials does so because a caller asks. The cost of crews that work this way is on the court record. A DOJ complaint unsealed on September 18, 2025 alleges that one member of the group took part in at least 120 intrusions from May 2022 to

September 2025, and that victims paid at least \$115 million in ransom.<sup>41</sup>

---

## Access has a market price

When asking fails, adversaries pay. On May 14, 2025 Coinbase disclosed that an actor had paid contractors or employees in support roles outside the United States to pull customer data from systems they could already reach for their jobs.<sup>42</sup> The path is  $H_s \mid / N \mid / H_c \mid H$ . No password or private key was taken. The data was then used to deceive customers into sending funds.<sup>42</sup>

---

## \$311.2M

in cash paid by Coinbase during 2025 on an incident that began with bribed support staff, made up of voluntary customer reimbursements and direct legal costs.

Coinbase Form 10-K for 2025<sup>43</sup>

Coinbase booked \$307 million of expense in the second quarter of 2025.<sup>44</sup> The affected data covered under 1 percent of monthly transacting users.<sup>45</sup>

The price on the other side of the transaction is also on record. The Cyber Safety Review Board found that Lapsus\$ offered as much as \$20,000 a week for access at telecommunications providers to conduct SIM swaps.<sup>46</sup> A former AT&T store employee sentenced on September 8, 2026 received a cut, typically \$1,000 to \$2,000, for SIM swaps his co-conspirators used to take over victims' bank accounts and wire money out.<sup>47</sup> A carrier store is someone else's identity desk. When the SEC's X account was taken over on January 9, 2024, the path began with a fraudulent ID card presented at a carrier store,  $H_t \mid / I / C$ .<sup>48,49</sup>

---


## Hiring became a way to issue credentials

Every new hire receives an account, a laptop and a place on the network. North Korea's remote IT worker schemes used that fact. They did not break into the companies. They applied.

Case 3 US companies employing North Korean IT workers via Christina Chapman laptop farm

2023 · Technology · US · North Korean IT workers with US facilitator Christina Chapman, state · Theft

I<sub>t</sub> | / H / P / V



Workers obtained jobs at 309 US companies using stolen and purchased identities of US nationals and netted more than \$17.1 million.

Confidence: high. Discovered by law enforcement.  
Dwell time not stated in sources.

Christina Chapman, sentenced on July 24, 2025 to 102 months, helped North Korean workers get jobs at 309 US companies and two foreign ones using 68 identities stolen from Americans. More than \$17 million moved through the scheme, and agents seized more than 90 laptops from her home.<sup>50</sup> Case 3 runs from stolen identities to deceived employers, company laptops in a house in Arizona and wages paid out. The domain that carries the intrusion is physical, a laptop on a shelf with a US address.

Chapman was not an outlier. On June 30, 2025 the Justice Department announced searches of 29 known or suspected laptop farms across 16 states, the seizure of 29 financial accounts and 21 websites, and about 200 computers.<sup>51</sup> Treasury says the schemes generated "nearly \$800 million in 2024." It gives no method for that number.<sup>52</sup>

The hiring checks were passed with real documents. A Maryland man passed an online interview with a company's


chief executive using his own driver's license and US passport, and was authorized for a federal PIV card while a worker in China did the job.<sup>53</sup> Three Georgia defendants helped overseas workers pass video interviews, drug tests and fingerprinting.<sup>54</sup> KnowBe4 hired a North Korean worker after four video interviews and clean background and reference checks. Its endpoint tooling caught him when it flagged malware loading on his new laptop on July 15, 2024.<sup>55</sup>

We read this as a seam between three owners that each did their part correctly. HR verified a person. IT shipped a device to an address. Identity issued an account to the person HR named. Nobody owned the question of who was actually at the keyboard.

# Most losses reported to the FBI ran through a person

The FBI's Internet Crime Complaint Center received 1,008,597 complaints reporting \$20.877 billion in losses in 2025. Business email compromise accounted for 24,768 complaints and \$3.05 billion, 14.6 percent of the total. Ransomware complaints reported \$32.3 million.<sup>56</sup> By our calculation from the IC3 tables, reported BEC losses were 94 times reported ransomware losses, and the six schemes in Figure 9 in which a person is deceived reported \$15.0 billion, 72 percent of all losses.

Figure 9 Reported BEC losses were 94 times reported ransomware losses in 2025



n is the number of 2025 complaints across the crime types shown. IC3 ransomware losses exclude downtime, remediation and lost business, and many ransomware victims do not file with IC3, so the ratio overstates the gap in total cost. Investment, BEC, tech support, government impersonation, employment and SIM swap are all schemes where a person is deceived.

Source: FBI IC3 2025 Internet Crime Report, three-year loss and complaint comparison tables, retrieved 2026-09-28. Bendings analysis. n = 207,240


The ratio needs its caveat next to it. IC3 ransomware losses leave out downtime, recovery and lost business, and many ransomware victims never file with IC3. Coinbase alone paid nearly ten times the IC3 ransomware total in cash on one incident. The ratio shows where reported losses concentrate, not what ransomware costs.

Gulf figures are thinner and point the same way. The UAE Financial Intelligence Unit estimated AED 1.24 billion in fraud losses for 2021 to 2023, from fund-recall data of 41 financial institutions. In its sample of 879 suspicious trans-

action reports from July 2023 to June 2024, vishing and phishing made up 16 percent, impersonation 10 percent and BEC 7 percent. It found that BEC perpetrators "relied heavily on social engineering techniques more than technical hacking."<sup>57</sup>

The reported role of AI is small. Only \$30.3 million of 2025 BEC losses, about 1 percent, came from complaints with AI-related information.<sup>56</sup> That is a floor, because the tag records only what victims knew.

Figure 10 About 1% of 2025 BEC losses carried a reported AI link



The AI tag depends on what victims knew and reported, so it is a lower bound; IC3 says many victims do not realize AI was used. Across all crime types IC3 received 22,364 complaints with AI-related information and \$893 million in losses in 2025.

Source: FBI IC3 2025 Internet Crime Report, 'AI References by Complaint Loss' and three-year loss comparison, retrieved 2026-09-28. Bendings analysis. n = 24,768

The best-known deepfake payment fraud, the fake video call with senior managers at Arup early in 2024, needed no system compromise. "None of our systems were compromised," Arup's chief information officer said.<sup>58</sup> A synthetic voice or face makes the request more convincing, and the control that failed was still the payment approval. Chapter 8 returns to AI.

## Where the data disagrees

Two findings cut against the argument of this chapter, and a third limits it.

First, most people-driven fraud has a single bend. BEC and impostor fraud usually run from a deceived person straight to a payment. The IC3 losses in Figure 9 are large, but they mostly cross one seam, between a person and a payment system. They support the claim that people are controls. They do not show intrusions bending through them.

Second, SIM swap losses reported to IC3 fell from \$72.7 million on 2,026 complaints in 2022 to \$17.4 million on 971 complaints in 2025, a drop of 76 percent by our calculation.<sup>59,56</sup> Losses from accounts taken over after a SIM swap may be filed under other crime types, and the fall may reflect reporting as well as fewer attacks. Still, the carrier seam, one of the most cited human-to-identity bends of the early 2020s, is not growing in this data.

Finally, the strongest numbers on help desk intrusions come from one vendor's caseload and from allegations in criminal complaints. Filings from victims rarely say what the help desk did. MGM Resorts' filing on its 2023 incident, for example, does not name the entry method, and we have not coded one.<sup>60</sup> Marks & Spencer's chairman told Parliament only that the entry "involved a third party". The help desk detail in that case rests on a press report.<sup>5,6</sup>

## What to do

- **Head of identity:** bring every process that resets a password, enrolls or removes an MFA device, or recovers an account under the identity program, including outsourced service desks and contact centers. Require a second channel that the caller does not control, and log each reset so it can be joined to the sign-ins that follow it.
- **Head of offensive security:** put pretext calls to the internal help desk, the outsourced service desk and the customer call center in every red team scope, using test identities set up for the exercise. Test enrollment and recovery for passkeys as well as passwords.
- **Procurement and supplier risk leads:** for every support contractor whose staff can reset credentials or read customer records, write into the contract the right to test that desk and to receive its logs, and require alerts on bulk record lookups by support staff.
- **Chief human resources officer, with the CISO:** treat onboarding as an access grant. Confirm that the person who interviewed is the person who uses the laptop, ship devices only to verified addresses, and limit what a new hire can reach until those checks pass.
- **Head of payments operations:** test the change of supplier bank details, the urgent executive request and the inbound customer instruction, with and without synthetic voice or video, and rehearse funds recall. In 2025 the IC3 Recovery Asset Team froze \$679 million of \$1.16 billion across 3,900 incidents reported quickly.<sup>56</sup>



5

Supplier seams are  
crossed often, tested rarely

Counting each supplier's breach once, about one in five incidents US companies disclosed came through a supplier, as did one in three intrusions in our corpus. Of 41 US and Gulf testing regimes, 9 put supplier-run systems in scope.

|                                                                       |    |
|-----------------------------------------------------------------------|----|
| Disclosures count filers, not supplier breaches                       | 34 |
| After May 2024, most incidents were no longer first filed as material | 35 |
| Six forms the supplier seam takes                                     | 35 |
| Emergency fixes cross the seam too                                    | 38 |
| Why almost nobody tests this seam                                     | 38 |
| Where the data disagrees                                              | 38 |

Between June 21 and June 26, 2024, seven listed US vehicle retailers filed their own Form 8-K reports with the Securities and Exchange Commission about the same event: an intrusion at CDK Global, the company whose software their dealerships used.<sup>61</sup> None of the seven reported an attack on its own systems. Each disclosed anyway, because part of its business ran on someone else's.

## Disclosures count filers, not supplier breaches

Item 1.05 of Form 8-K has required US public companies to disclose material cybersecurity incidents since December 18, 2023, and smaller reporting companies since June 15, 2024.<sup>62,63</sup> We collected every Item 1.05 filing to September 28, 2026, and every 8-K we could find that disclosed an incident under Items 8.01 or 7.01 instead. They describe 135 incidents at 130 companies. We coded each first disclosure and checked the rules against a second, independent coding and a hand-read sample.<sup>61</sup>

## 36 of 135

companies disclosing a cybersecurity incident on Form 8-K from December 18, 2023 to September 28, 2026 said it occurred at, or came through, a third party. Counting each supplier's breach once, at most 28 of 127.

SEC EDGAR, Bendings analysis<sup>61</sup>

In 36 of the 135 first disclosures, 27 percent, the filer said the incident occurred at or came through a third party. These are filers, not supplier events. Seven filings concern one CDK Global incident, two concern Change Healthcare and two concern Evertec's 2026 incident. Grouping filers whose filings name the same supplier leaves 28 supplier events among 127 incidents counted that way, 22 percent. Twenty-one of the 28 do not name the supplier and may overlap, so 28 is a ceiling.

Figure 11 36 of 135 incidents disclosed on Form 8-K occurred at or came through a third party



Coded from the first filing's text: the incident happened at, or came through, a vendor, supplier, service or software provider, business partner or processor. Hiring outside experts does not count. Counts: Dec 2023 to 2024: 17 of 58; 2025: 7 of 36; 2026 to Sep 28: 12 of 41; All: 36 of 135. Later disclosures may reveal a third-party path the first 8-K did not state.

Source: SEC EDGAR Form 8-K and 8-K/A filings, 2023-12-18 to 2026-09-28 (full-text search and submissions API), retrieved 2026-09-28. Bendings analysis. n = 135

Among financial firms, meaning banks, insurers, lenders, securities firms and other companies with SIC codes 6000 to 6799, it was 10 of 28 incidents, against 26 of 107 in other sectors, a difference chance alone could produce at this size. Third party is also the code our two codings agreed on least, a kappa of 0.70 across 142 filings, and we settled each of the 17 disagreements by reading the filing.


One filer found the seam one step further out. The Federal Home Loan Bank of New York told the SEC in March 2024

that its incident came from "a fourth-party vendor (i.e., a vendor of a Bank vendor)."<sup>64</sup> A bank's supplier program stops at its contracts. Its exposure does not.

# After May 2024, most incidents were no longer first filed as material

On May 21, 2024, SEC staff said that incidents not yet judged material could be disclosed under Item 8.01, "Other Events", rather than Item 1.05.<sup>65</sup> Before that statement, 17 of 23 new incidents were first disclosed under Item 1.05. After it, 33 of 112 were. A law firm tracker that follows these filings reports the same shift.<sup>66</sup>

Figure 12 After the SEC staff statement of May 2024, 33 of 112 new incident disclosures used Item 1.05; before it, 17 of 23 did



One bar segment per incident, placed in the quarter of its first 8-K. 2023-Q4 covers 2023-12-18 to 2023-12-31 and 2026-Q3 covers 2026-07-01 to 2026-09-28. Incidents first disclosed before 2023-12-18 are excluded. The staff statement (2024-05-21) asked companies to use Item 8.01 for incidents not yet determined material.

Source: SEC EDGAR Form 8-K and 8-K/A filings, 2023-12-18 to 2026-09-28 (full-text search and submissions API), retrieved 2026-09-28. Bendings analysis. n = 135

Counted by filer, third-party incidents lean to the "other events" side, 27 of the 85 first disclosed under Items 8.01 or 7.01 against 9 of the 50 first disclosed under Item 1.05. The seven CDK filers make most of that gap. Counting each supplier event once, it is 19 of 77 against 9 of 50, with-in chance.

## Six forms the supplier seam takes

The corpus lets us say more than "third party". Of its 78 documented intrusions, 26 pass through a supplier, in a small number of patterns, each with a different control on the victim's side and a different owner on the other.

| Form                               | What crosses                                                                | Examples in the corpus                                                                                              |
|------------------------------------|-----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Supplier credentials               | A contractor's or other supplier's account used on the victim's own systems | Target 2013, Home Depot 2014, OPM 2014 <sup>8,67,68</sup>                                                           |
| Outsourced people                  | A supplier's help desk or support staff deceived or bribed                  | Okta via Sitel 2022, Caesars 2023, Marks & Spencer 2025, Coinbase 2025 <sup>69,32,6,42</sup>                        |
| Software and build                 | A poisoned update, package or build tool that the victim runs               | NotPetya via M.E.Doc 2017, SolarWinds 2020, 3CX 2023, the European Commission via Trivy 2026 <sup>70,71,72,73</sup> |
| Remote access and management tools | A supplier's tool or key that already has a path into the victim            | Kaseya VSA 2021, US Treasury via BeyondTrust 2024 <sup>74,75</sup>                                                  |
| SaaS integrations and tenants      | Tokens or data held by one cloud service inside another                     | Snowflake tenants 2024, Salesloft Drift 2025, 8x8 via Klue 2026 <sup>76,77,78</sup>                                 |
| Processors and platforms           | A payment processor or platform provider whose failure is the victim's loss | RAKBANK and Bank Muscat 2013, Storm-0558 2023, Bybit via Safe{Wallet} 2025 <sup>15,16,79,1</sup>                    |


In our view, a test run with the supplier taking part can walk supplier credentials, outsourced people, remote access tools and SaaS integrations, because each crosses a connection both sides can see. Processors and platforms can be walked as far as the provider opens its side. A poisoned build cannot be walked from the customer's side at all. Only the supplier's own testing reaches it.

In the European Commission case of 2026, attackers compromised the build environment of Trivy, a widely used open-source security scanner, and published poisoned releases that gave up a Commission cloud key when run.<sup>73,80</sup> A tool bought to find weaknesses was the way in, as chapter 3 found for the edge.

Case 4 European Commission Europa web hosting cloud account (Trivy supply chain)

2026 · Government · Europe · TeamPCP (supply-chain compromise); data published by ShinyHunters, criminal · Extortion

N<sub>s</sub> / I<sub>s</sub> / N<sub>s</sub> | N / C . C



- 1

N

Exploited a misconfiguration in the automated build environment of the Trivy open-source security scanner to extract a privileged access token<sup>80</sup> supplier
- 2

I

Kept access through credentials that the project's incomplete credential rotation left valid<sup>80</sup> supplier
- 3

N

Published a malicious Trivy release and redirected trusted version references for its build actions to malicious code<sup>80</sup> supplier
- 4

N

The compromised scanner, received through normal update channels, ran in the Commission's environment and exposed a cloud access key with management rights<sup>73</sup>
- 5

C

Used the stolen cloud key, attached a new access key to an existing user to avoid detection, and carried out reconnaissance<sup>73</sup>
- 6

C

Exfiltrated website data and outbound email records from the cloud account, which were later published on a leak site<sup>73,81</sup>

About 91.7 GB (compressed) of data was taken from an AWS account behind the europa.eu web hosting service, including names, email addresses and email content, affecting up to 71 hosting clients: 42 Commission services and at least 29 other Union entities. ShinyHunters published the dataset on March 28, 2026; no websites were taken offline.

Confidence: high. Discovered by the victim. Dwell time 5 days.

A case too recent for our corpus shows the SaaS form at machine speed. In October 2026 Microsoft attributed a compromise of ConnectWise's cloud backup applications to Storm-0940, a group it calls China-state aligned. Holding the identities of several ConnectWise cloud applications that customers had already granted access, the group's AI agents tried to sign in to more than 2,250 Microsoft 365 tenants. In some they went on to change application credentials and privileged roles, which Microsoft warns can stay in place as a backdoor after the application itself is fixed. It adds that an affected application, or even a successful sign-in, does not show that data was taken.<sup>82</sup>

Harm runs back the other way. In 48 of the 78 corpus cases it reached an owner other than the victim, most often customers. Kaseya's compromised management tool reached managed service providers and, through them, their customers.<sup>74</sup> When Change Healthcare cut its systems off in 2024, claims and payment flows across US health care were disrupted with them.<sup>17</sup> Reconnecting is the harder half. In the Bank of England's 2024 cyber stress test, firms said the time needed to get a third party's assurance that it was safe to reconnect might exceed their own impact tolerances.<sup>83</sup>

## Emergency fixes cross the seam too

The PHP flaws fixed in October 2026, described in chapter 8, sat in the runtime under a website, and many affected organizations do not run that runtime themselves. Cloudflare told customers that updating a site's code or plugins would not fix them, and that those whose host manages PHP should ask the host.<sup>84</sup> The Dutch National Cyber Security Centre gave the same advice about hosting providers and IT suppliers. A UK minister told the House of Lords that small businesses would most probably be relying on their managed service providers.<sup>85,86</sup>

The fix also traveled through suppliers before release. PHP sent it in advance to Linux distributions and long-term support vendors so their users could update the moment it was disclosed.<sup>87,88</sup> By the PHP Foundation's account, one distributor published patched images more than six hours early, letting anyone study the fix while every other PHP user waited.<sup>88</sup> An embargo crosses as many owner seams as it has recipients.

## Why almost nobody tests this seam

The victim cannot authorize a test of its supplier's systems or people, and the supplier rarely invites one. So the seam is examined by questionnaire and contract, which are assurance, not tests. In our view, none of the six forms would show up in a questionnaire answer, because each exploits the connection rather than a gap either side would list.

The testing rules reflect this. Chapter 9 found that 9 of 41 US and Gulf testing regimes put systems a supplier runs for the regulated firm inside the test, 7 of 12 in the US and 2 of 29 in the Gulf, where 10 rely on supplier assurance alone. The EU's DORA is the only text we found that gives a tester authority to cross the seam no single customer can authorize. When several financial firms share an ICT provider they may run a pooled test, and each pooled test must include at least one scenario on the provider's own systems.<sup>89,90</sup>

## Where the data disagrees

Most disclosed incidents do not involve a third party. In 99 of the 135 SEC incidents, the first filing states no third-party path. A filing may leave out a supplier's role, but we can only code what it says. The corpus has the opposite bias to the SEC count, since it codes a campaign such as SolarWinds or Snowflake once, however many victims it had. In nine corpus paths the first step at a supplier or customer is inferred, not documented (chapter 2).

Neither dataset shows a trend. The corpus is too small to show one (chapter 2), and by filer the SEC share was 17 of 58 from December 2023 through 2024, 7 of 36 in 2025 and 12 of 41 in 2026 to September 28.

## What to do

- **Supplier risk lead:** list every supplier whose people, credentials, software or integrations reach your systems, with the form of seam each one opens, who owns the connection on both sides and who applies emergency fixes to what the supplier runs for you.
- **Procurement:** write into renewals the right to run a joint scenario across the connection and access to the supplier's relevant logs for the test window. Require notice of an incident in time for you to meet your own first deadline. For a bank licensed in Bahrain, Qatar or Kuwait that can be one hour, in the UAE four hours (chapter 10).<sup>91,92,93,94</sup> Extend the notice duty to the supplier's own critical suppliers.
- **Head of cloud and SaaS:** inventory third-party integrations and their token scopes in every tenant, remove unused ones, and alert on bulk data export by an integration.
- **Head of offensive security:** run the supplier-access scenario from chapter 12 with one critical supplier taking part, and report the paths the test could not follow because authority stopped.
- **Regulators:** in outsourcing and cloud rules, require contracts with shared providers to include taking part in a pooled test led by their regulated customers, with at least one scenario on the provider's own systems, as DORA does. A central bank can then coordinate the pool without supervising the provider.



# 6

## Few tests cross the door, camera or valve

The network ends at a door, a camera, a cash machine and a valve. Network monitoring and most testing rules stop short of each of these seams, and adversaries have crossed all four.

|                                                            |    |
|------------------------------------------------------------|----|
| Physical steps are rare in the public record               | 40 |
| Cameras and door controllers are on the exploited list     | 41 |
| The cash machine is a computer with a front door           | 42 |
| OT attacks moved from defacement to disabling safety logic | 42 |
| IT failures stopped OT as often as adversaries did         | 44 |
| Where the data disagrees                                   | 44 |

On February 19, 2026, the FBI told US banks it knew of 1,900 ATM jackpotting incidents reported since 2020, more than 700 of them in 2025, with more than \$20 million lost that year. The crews opened the machines most often with widely available generic keys, then pulled the hard drive, loaded malware onto it and put it back.<sup>95</sup>

The FBI's advice is about what the bank's network cannot see. It asks banks to audit removable storage use and compare file hashes on the machine against a known image, checks that together catch "physical intrusion and malware staging events that would otherwise evade network-based monitoring".<sup>95</sup> The attack starts at a lock and ends in a cash dispenser. Nothing in between crosses the network the bank watches.

Testing rules rarely reach the building or the plant. Of 41 US and Gulf testing regimes we coded, 1 requires a physical intrusion test. It is Saudi Arabia's controls for operational technology, OT from here on, and it applies only at the two highest OT facility levels.<sup>96</sup> OT is not addressed in 31 of the 41. Kuwait's central bank is the exception worth copying. Its 2025 framework requires annual penetration tests of IoT and OT and names CCTV and building management systems as examples.<sup>93</sup> Qatar's penetration testing accreditation standard excludes ICS and OT testing from what it accredits.<sup>97</sup> Chapter 9 sets out the regimes in full.

## Physical steps are rare in the public record


Of the 78 documented paths in our corpus, 4 include a physical step, coded P in bend notation. The GRU's 2016 operation against the Canadian Centre for Ethics in Sport started with two officers carrying close-access equipment to a hotel in Lausanne and compromising its Wi-Fi, **P<sub>t</sub> | / N / I / N**.<sup>98</sup> In the DarkVishnya cases of 2017 and 2018, visitors posing as couriers or job seekers plugged small devices into the networks of at least eight Eastern European banks, **H / P / I / N / V**.<sup>99</sup> The third is the US lap-top farm in chapter 4.<sup>50</sup>

The fourth, drawn in full below, shows how short the physical step can be.

Case 5 **Organization A (Volexity customer) via nearby Wi-Fi networks**

2024 · Other · Other · GruesomeLarch (APT28), state · Espionage

I | / N<sub>t</sub> | / P / N



- 1 I Validated employee credentials by password-spraying a public-facing service protected by MFA<sup>100</sup>
- 2 N Breached nearby organizations (B and C), including through a VPN without MFA, and found dual-homed systems with Wi-Fi adapters<sup>100</sup>  
third party

- 3 P Used the neighbor's Wi-Fi adapter within radio range to join Organization A's enterprise Wi-Fi, which required only username and password<sup>100</sup>
- 4 N Moved inside Organization A and attempted to dump the Active Directory database via a volume shadow copy<sup>100</sup>

APT28 reached Organization A's enterprise Wi-Fi from the networks of neighboring organizations to collect data from people working on Ukraine. Volexity observed an attempt to steal the Active Directory database.

Confidence: high. Discovered by a third party. Dwell time not stated in sources.

In February 2022 APT28 held valid passwords for Organization A, but its public service asked for MFA and its enterprise Wi-Fi did not. The operators breached two neighboring organizations, found machines there with both a wired connection and a Wi-Fi adapter, and joined Organization A's wireless network from next door.<sup>100</sup> The path is I | / N<sub>t</sub> | / P / N. The physical step was radio range, borrowed from someone else's building.

The GRU has used close access more than once. In April 2018 the Dutch MIVD disrupted a GRU Wi-Fi operation run from a car parked beside the OPCW in The Hague, and the Netherlands expelled four officers.<sup>101</sup> The corpus's four cases in thirteen years count only what was explained in public.

## Cameras and door controllers are on the exploited list

Inside the building, the physical security system is itself a network of small computers. Our classification of CISA's Known Exploited Vulnerabilities catalog, version 2026.09.27, finds 22 of 1,728 entries, 1.3 percent, in cameras, network or digital video recorders, and access-control or time-attendance systems. Seventeen of the 22 were added from January 2024 onward.<sup>20</sup>

**Figure 13** 17 of the 22 camera, video recorder and access-control flaws in KEV were added from 2024 onward



22 of 1728 KEV entries (1.3%) are cameras, network or digital video recorders, or access-control and time-attendance systems, found by keyword screen and hand review; the method may undercount. 2021 covers November and December only (the catalog began 2021-11-03) and 2026 runs to 2026-09-27. Building management and industrial controllers are excluded.

Source: CISA Known Exploited Vulnerabilities catalog, version 2026.09.27 (1728 entries), retrieved 2026-09-28. Bendings classification and analysis. n = 22

The share is small. The risk lies in what watches these devices. In one case S-RM investigated, endpoint protection quarantined Akira's ransomware on the victim's hosts. The operators found an unsecured webcam that ran no such protection and deployed the ransomware from it, encrypting files across the network.<sup>102</sup> Our reading of this one responder's case is that a device the security team does not watch is a device the adversary can work from.

Cameras also look outward. US and allied agencies reported in May 2025 that GRU unit 26165 had targeted internet cameras to watch shipments of aid to Ukraine.<sup>103</sup> Check Point reports a spike from February 28, 2026 in Iran-linked scanning and exploitation attempts against Hikvision and Dahua cameras in Qatar, Bahrain, Kuwait and the UAE. It says the activity is consistent with Iran using compromised cameras for battle damage assessment, but it does not report confirmed compromises and gives no counts in its text.<sup>104</sup> On March 5, 2026, CISA added a Hikvision flaw from 2017, one of those Check Point lists, to the catalog.<sup>105</sup>

## The cash machine is a computer with a front door

The FBI's jackpotting count covers only what was reported to it.<sup>95</sup> In Nebraska, federal prosecutors had charged 87 defendants in one jackpotting conspiracy by January 26, 2026.<sup>106</sup>

The ATM also reaches inward. Group-IB describes an intrusion at an unnamed bank in which UNC2891 placed a

Raspberry Pi with a 4G modem on the same network switch as an ATM and used it for remote access to the bank's network.<sup>107</sup> In bend notation that path starts **P / N**, at the switch beside the machine.

Our reading is that three owners meet at the machine. Branch facilities hold the lock, a supplier may service the machine, and the security operations center watches the network. The control the FBI recommends, image integrity on the machine, belongs to none of them by default.

## OT attacks moved from defacement to disabling safety logic


Six of the 78 paths in the corpus reach OT. In December 2015 intruders operated breakers remotely at three Ukrainian distribution companies and cut power to about 225,000 customers.<sup>108</sup> A year later Industroyer spoke four industrial protocols directly to switches and breakers.<sup>109</sup>

In 2017 the target moved to the system meant to stop a plant from hurting people. Malicious code on the Triconex safety controllers at a petrochemical facility tripped one controller on June 2 and six on August 4. INL's case study counts 17 days of halted operations across the two trips and about \$938,000 in lost revenue, with no physical damage.<sup>110</sup> US official sources place the facility in the Middle East without naming the country.<sup>111</sup>

Case 6 Triton safety-system attack on a Saudi petrochemical facility

2017 · Energy · Saudi Arabia · TEMP.Veles / Central Scientific Research Institute of Chemistry and Mechanics (TsNIIKhM), state · Sabotage

E / N / I / O



Malicious code on the plant's Triconex safety controllers tripped one controller on June 2, 2017 and six controllers on August 4, 2017, shutting the facility down twice; INL's CyOTE case study reports about 10 days of lost operations after the second trip. No physical damage occurred.

Confidence: high. Discovered by the victim. Dwell time not stated in sources.

The path is E / N / I / O. Every step sat on the victim's own estate. The seam was between the IT team and the engineers who own the safety controllers.

The US record since then follows the same arc with simpler methods. From November 2023 to January 2024, IRGC-linked actors calling themselves CyberAv3ngers compromised at least 75 Unitronics controllers, at least 34 in US water systems, through default or absent passwords, and left a defacement image.<sup>115</sup> Advisory AA26-097A, first issued April 7, 2026 and updated July 22, reports that Iranian-affiliated actors since at least March 2026 have altered logic on internet-exposed controllers in water, energy and government facilities. The changes "disabled critical shutdown and alarm logic, allowing systems to enter unsafe conditions without notifying operators", and some victims suffered operational disruption and financial loss.<sup>19</sup>

Both advisories tie the actors to Iran. The 2023 intrusions changed what a screen showed. The 2026 intrusions changed what the controllers would do in an emergency.

Poland shows the other direction of travel, from the edge inward. On December 29, 2025, more than 30 wind and solar farms lost remote monitoring and control from the grid operator after attackers came in through firewall and VPN devices at the connection points and damaged substation control devices. Generation continued.<sup>116</sup>

## IT failures stopped OT as often as adversaries did

### 7 vs 6

In 7 of 78 documented intrusions, harm reached OT only because the victim or someone downstream shut it down or could not run it. In 6, the adversary acted in OT itself.

Bendings incident corpus, 2013 to September 2026

The corpus codes consequences separately from the adversary's path and counts harm by decision only in domains the adversary never acted in, so the two groups in the callout cannot overlap.

Chapter 2 lists the best-known shutdowns, among them Colonial Pipeline and Norsk Hydro. Wind farms in central Europe were disrupted when the Viasat satellite modems they relied on were wiped, although no one touched their controls.<sup>117</sup> Nevada's restoration list after its 2025 ransomware incident included automatic locks at state facilities.<sup>118</sup>

We read this as the most useful OT finding in the chapter. A victim's own decision to isolate or halt plant, because IT can no longer be trusted, reached OT about as often as adversaries did. That decision has an owner, a trigger and a cost, and it can be tested without touching a controller.

## Where the data disagrees

We found no official US or GCC case of a network implant planted in a bank branch. The best-documented cases come from vendors, Kaspersky's DarkVishnya work in Eastern Europe and Group-IB's unnamed bank. Jackpotting is well documented, but it attacks the device and does not open a way into the bank's network.

Some hardware implant fears have not been confirmed. A House committee letter in March 2024 reported unexplained cellular modems on port crane components, and we found no executive-branch forensic finding behind it.<sup>119</sup> A press report says a US Department of Energy review of about 30 Chinese-made inverters found no definitive evidence of malicious wireless functions, and two cases of undocumented communications judged not malicious.<sup>120</sup> We could not open the review itself.

Physical steps are rare in the public record, 4 of 78, and that may say more about disclosure than about attacks. Our reading is that victims rarely publish how someone got into a building, and that a path starting with a device on the network tends to be written up as starting on the network. The FBI's jackpotting count shows that physical attacks on bank devices are common. It cannot tell us how often a physical step into a network goes unreported.

The OT comparison is 7 cases against 6 in a sample we chose, too few to rank shutdown decisions above adversary action.

Finally, the most serious OT cases here are state operations against energy and water. None of the six OT paths is in finance, and only Triton is in the Gulf. For a bank, the physical and OT seams are the branch, the ATM, the camera and the building systems.

## What to do

- **Heads of physical security:** put every camera, video recorder and door controller on the asset register the security operations center monitors, and check the list against the KEV catalog each month.
- **Heads of offensive security:** when a physical test gets in the door, let it continue to the network. Plant a device, join the Wi-Fi from the lobby or car park, and report how long it ran before anyone noticed.
- **Heads of ATM and branch operations:** adopt the FBI's hard drive image and hash checks, replace the standard locks on ATM fascias, and agree in writing who owns an alert from a machine.
- **OT and plant engineering leads:** confirm that no controller or safety system is reachable from the internet, change default passwords, and test that shutdown and alarm logic still fires after any remote change.
- **Operations and resilience leads:** write down who decides to isolate or halt OT when IT is compromised, on what trigger and at what daily cost, then rehearse it.
- **SAMA and the CBK, for tests under FEER and CORF:** our proposal is to publish the terms for a physical scenario: letters of authority, notice to police, building owners and, near critical infrastructure, the state security authority, and the handling of staff data. In the

Gulf in 2026, a physical test no security authority knows about puts the tester at risk as well as the test.

- **Qatar's National Cyber Security Agency:** our proposal is to bring ICS and OT testing into its penetration testing accreditation.



# 7

## In the Gulf, a cloud region has a street address

In 2026 drone strikes took down cloud facilities in the UAE and Bahrain, and some data held only there was lost for good. Gulf rules that keep data and testers in the country now sit against recovery plans that assume another region.

|                                                                                           |    |
|-------------------------------------------------------------------------------------------|----|
| Two cloud regions failed together, and one lost data for good                             | 47 |
| Residency rules and recovery advice point in opposite directions                          | 47 |
| Leak sites name the Gulf somewhat more often than in 2022                                 | 48 |
| Three Gulf central banks want a first report within the hour, and the public hears little | 50 |
| Interference is a standing condition                                                      | 52 |
| Where the data disagrees                                                                  | 52 |

At about 4:30 in the morning Pacific time on March 1, 2026, objects struck an Amazon Web Services data center in the United Arab Emirates, "creating sparks and fire." The fire department cut power to the building and its generators. The facility held mec1-az2, one of three availability zones in AWS's ME-CENTRAL-1 region. On September 15, AWS said it could not restore any resources or data held only in that zone, or held only in its Bahrain region.<sup>121</sup>

Physical attacks on Gulf infrastructure were not new. Saudi Aramco reported attacks with projectiles on its Abqaiq and Khurais plants in September 2019, and in January 2022 the Security Council condemned attacks in Abu Dhabi that, in its account, the Houthis committed and claimed.<sup>122,123</sup> In 2026 the physical seam reached the cloud. A cloud region is a set of buildings in one country, and in March some of them were struck.

## Two cloud regions failed together, and one lost data for good

AWS's public health dashboard is the primary record. On March 2 it said that "both affected regions have experienced physical impacts to infrastructure as a result of drone strikes." Two UAE facilities were struck directly, and a strike close to one Bahrain facility caused physical damage. Two of the three UAE zones were "significantly impaired", and AWS told customers to "recover from remote backups into alternate AWS Regions, ideally in Europe." On April 30 it said the Bahrain region was "currently unavailable." Its September 15 update for Bahrain reads:<sup>121</sup>



## The damage to our infrastructure spanned multiple Availability Zones and exceeded what our regional and multi-AZ services are designed to withstand.

AWS Health Dashboard, ME-SOUTH-1 (Bahrain), September 15, 2026<sup>121</sup>

AWS does not say who launched the drones, and we found no official attribution of the strikes on these facilities. They came in a war that, in the UN's account, began on February 28, 2026, with Israeli and US airstrikes on Iran. On March 11 the Security Council adopted resolution 2817, which condemned "in the strongest terms" Iran's attacks against

Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, the United Arab Emirates and Jordan. The vote was 13 to none, with China and Russia abstaining.<sup>124</sup> On July 2 the Council met in emergency session after the United States and Iran resumed fighting in late June, despite a peace deal reached between them. The UN's coverage of that session says attacks on regional infrastructure had continued.<sup>125</sup>

In bend notation the customer's loss reads  $P_s / C_s$ , a physical event at a supplier's building that took down the supplier's cloud service. No adversary touched the customer's systems, yet the customer bore the loss, which is what chapter 2 means by consequences that bend. AWS describes S3, its storage service, as "designed to withstand the total loss of a single Availability Zone." Our reading is that strikes across a country are not the single-building failure that design assumes. AWS says most customers had moved their workloads before the second Bahrain zone failed in April. Those who had not lost whatever was held only there.<sup>121</sup>

## Residency rules and recovery advice point in opposite directions

AWS recommended regions "in the United States, Europe, or Asia Pacific, as appropriate for your latency and data residency requirements."<sup>121</sup> For many regulated firms in the Gulf, those requirements point the other way.

- **United Arab Emirates.** The Central Bank of the UAE's Operational Risk Management Regulation, in force since September 14, 2026, says that "an LFI's Master System of Record must be continuously maintained and stored within the UAE, including in the case of Outsourcing." An LFI is a licensed financial institution. Branches of foreign banks may instead keep an up-to-date copy in the UAE, with the central bank's approval (Art. 8.9).<sup>126</sup>
- **Bahrain.** The Central Bank of Bahrain requires prior approval for outsourcing outside Bahrain but excludes

"cloud data services" from that requirement.<sup>127</sup>

- **Qatar.** The Qatar Central Bank's Cloud Computing Regulation says that "an Entity must ensure that PII and financial information is processed within Qatar only" (Art. 21.4).<sup>128</sup> Its Data Handling and Protection Regulation allows a financial institution's "primary environment" outside Qatar "only after receiving explicit approval from QCB" (7.7).<sup>129</sup>
- **Saudi Arabia.** The National Cybersecurity Authority's (NCA) controls require critical systems to be hosted inside the organization or with government bodies or Saudi companies that comply with its cloud controls (CSCC 4-2-1-1).<sup>130</sup> Its 2024 Essential Cybersecurity Controls sent data localization to the National Data Management Office, whose rules we did not review.<sup>131</sup>

The same rules reach the people who test and watch. The NCA prohibits remote access to critical systems from outside Saudi Arabia (CSCC 2-2-1-1).<sup>130</sup> Managed security operations centers that use remote access "shall be fully located in the Kingdom" (ECC 4-1-3-2).<sup>131</sup> The red team framework of the Saudi Central Bank (SAMA) says data a tester exfiltrates should be stored "locally (not at cloud providers)."<sup>132</sup>

For a UAE bank, the advice to restore in Europe now meets a rule that its master record stay in the UAE. That rule took effect six months after the strikes, so it will govern the next loss. In Qatar, a primary environment abroad needs the regulator's explicit approval, and in a loss like this one it would be needed within days. Bahrain's cloud exclusion makes the tension there smaller. We have no data on whether any approval was sought. Keeping financial data in the country is a deliberate policy choice, and we do not read these rules as wrong. Our view is that a residency rule should say in advance where data may go if every facility in the country is lost. Otherwise the bank and the regulator decide it during the loss.


fice or market. The strict rule counts only posts with a Gulf domain, flag or labeled country field, and undercounts. From January 1 to September 28 of each year, the broad rule's Gulf share of all posts went from 8 of 2,733 (0.29 percent) in 2022 to 68 of 7,790 (0.87 percent) in 2026. In 2021 it was 0.38 percent.

Much of that rise comes from the tracker. The broad rule works mostly by reading post descriptions, and the share of posts that carry one went from 17 percent in 2022 to 74 percent in 2026. Among posts with a description, the broad Gulf share went from 3 of 475 (0.63 percent) to 62 of 5,737 (1.08 percent), and it fell in 2023 and again in 2025. The strict share was 0.07 percent in 2022, 0.33 percent in 2023 and 0.22 percent in 2026, on 2 to 17 posts a year. We read this as a modest rise in how often Gulf organizations are named, with no steady trend.

## Leak sites name the Gulf somewhat more often than in 2022

Our analysis of the RansomLook leak-site tracker assigns posts to the Gulf in two ways.<sup>133</sup> The broad rule counts posts whose title or description names a Gulf country, city or national domain, and overcounts. In a hand review of 60 posts found only by it, 12 were foreign firms with a Gulf of-

**Figure 14** Much of the rise in the Gulf share of leak-site posts follows the tracker recording more post descriptions




Country is assigned by Bendings, not by the tracker. Rule A (strict) counts a post with a GCC country-code hostname (.sa .ae .qa .kw .bh .om), a GCC flag emoji or a labeled country field in its title or description, and undercounts. Rule B (broad) adds any whole-word GCC country, demonym or major-city name, and overcounts. In a hand review of 60 posts found only by rule B, 47 were Gulf organizations and 12 were foreign firms that mention a Gulf office or market. Leak-site posts are criminal claims; victims who paid, were never posted or post as bare names without Gulf markers are missing. Denominator is every deduplicated RansomLook post in the window, whatever the victim country, or for the third line every post that carries a description; n is all posts summed over the six windows. Posts with a description rose from 2 percent of all posts in 2021 and 17 percent in 2022 to 74 percent in 2026, and rule B mostly reads descriptions, so much of the rise in its share of all posts reflects the tracker recording more text. Rule A also reads bare-domain post titles, which fell from 39 percent of posts in 2023 to 20 percent in 2026. The third line is blank for 2021, when only 24 posts carried a description.

Source: RansomLook API (ransomlook.io, CC BY 4.0), posts discovered 2021-01-01 to 2026-09-28, retrieved 2026-09-28, deduplicated by group and title. Bendings analysis. n = 25,069

A post is a criminal claim, not a confirmed incident, victims who paid quietly never appear, and the country assignment is ours. Under the broad rule, UAE organizations made up

82 of the 162 Gulf posts across the six windows, and 37 of 68 in 2026.

**Figure 15 UAE organizations were about half of Gulf organizations named in leak-site posts (82 of 162 posts)**



UAE firms may name "Dubai" or "UAE" in public descriptions more often than other Gulf firms. Read the split as a lead, not a risk ranking.

Country is assigned by Bendings, not by the tracker. Rule A (strict) counts a post with a GCC country-code hostname (.sa .ae .qa .kw .bh .om), a GCC flag emoji or a labeled country field in its title or description, and undercounts. Rule B (broad) adds any whole-word GCC country, demonym or major-city name, and overcounts. In a hand review of 60 posts found only by rule B, 47 were Gulf organizations and 12 were foreign firms that mention a Gulf office or market. Leak-site posts are criminal claims; victims who paid, were never posted or post as bare names without Gulf markers are missing. Rule B. Posts naming more than one GCC state are shown separately.

Source: RansomLook API (ransomlook.io, CC BY 4.0), posts discovered 2021-01-01 to 2026-09-28, retrieved 2026-09-28, deduplicated by group and title. Bendings analysis. n = 162

The only official ransomware figure we found from any of the six states comes from the UAE Cyber Security Council's 2025 report with CPX. It says "the number of overall ransomware groups active in the UAE increased by 58 percent in 2024," and its source for that chart is a leak-site tracker, ransomfeed.it.<sup>134</sup> The official figure and ours rest on the same kind of data, so neither confirms the other.

## Three Gulf central banks want a first report within the hour, and the public hears little

### 1 hour

First report to the central bank after a serious cyber incident in Bahrain, Qatar and Kuwait. US bank regulators allow 36 hours from the bank's determination.

CBB, QCB, CBK, Federal Reserve and OCC rules <sup>91,92,93,135,136</sup>

The UAE central bank's new rule sets four hours for a first notice of a significant operational risk event.<sup>94,137</sup> Gulf rules also test more often, and SAMA and the Central Bank of Kuwait require threat-led red teaming on a fixed cycle,

which none of the seven US financial regimes we reviewed does.<sup>132,93</sup> Chapters 9 and 10 set out the rules.

What reaches the regulator is shared inside the sector, not with the public. The NCA says it shares its alerts and reports "with the relevant entities through its designated communication channels."<sup>138</sup> The UAE Cyber Security Council announces campaigns through the state news agency in sector terms, with no victim named.<sup>139</sup> Kuwait's ministries name themselves through the state news agency and describe restoration, not the entry point.<sup>140</sup> One Gulf regulator publishes aggregate findings. The Dubai Financial Services Authority's Cyber Thematic Review 2024, built on a questionnaire to the firms it supervises, found their compliance with third-party risk management practices below 70 percent.<sup>141</sup> US disclosure is thin too. SEC rules let companies withhold technical detail, and only 5 of 135 companies' first Form 8-K incident disclosures in our analysis named the attacker or called it state-sponsored (chapter 5).<sup>62,61</sup>

The corpus reflects this, but its Gulf and US cases were built from different kinds of source and cannot be compared closely. In two Gulf cases, the Qatar News Agency in 2017 and Kuwait's Ministry of Health in 2024, the government that owned the victim published the account.<sup>142,140</sup> Aramco and QNB answered press questions.<sup>143,144</sup> No Gulf compa-

ny in the corpus published an account of how the attackers got in. Four of the 15 Gulf cases rest on press reports alone, and six do not disclose the first step.


Both Gulf bank losses in the corpus, at RAKBANK and Bank Muscat, are documented by US prosecutors. In the

RAKBANK case, limits on prepaid card accounts were removed inside an unnamed card processor, and cashing crews withdrew about \$5 million from ATMs in about 20 countries on December 22, 2012.<sup>15,145</sup> The path never touched the bank's own systems.

Case 7 RAKBANK prepaid card cash-out via card processor

2013 · Finance · United Arab Emirates · 'Unlimited Operations' cybercrime organization led by Ercan Findikoglu, criminal · Theft


?<sub>s</sub> / I<sub>s</sub> / V<sub>s</sub> | V<sub>t</sub>



On December 22, 2012 cashing crews made more than 4,500 ATM withdrawals in about 20 countries with compromised RAKBANK prepaid card data, causing about \$5 million in losses to the card processor and RAKBANK (DOJ).

Confidence: medium. Discovered by not stated. Dwell time not stated in sources.

Figure 16 Fifteen documented Gulf intrusions, six with no disclosed first step



Paths that start with a dashed lead-in are cases where the sources do not say how the adversary got in.

Source: Bendings incident corpus, coded from public primary sources, snapshot September 28, 2026. n = 15 cases

# Interference is a standing condition

GNSS interference recurs in parts of Gulf airspace and waters. In the European Union Aviation Safety Agency's 30-day ranking dated September 23, 2026, the Jeddah flight information region placed 10th and the Emirates region 15th of the 51 regions listed for GNSS interference.<sup>146</sup> At sea, the Joint Maritime Information Center reported on August 9, 2026, that "recurring GNSS interference remained consistent with prior reporting."<sup>147</sup>

IT tests rarely cover the seam between a satellite signal and a system that trusts it. Our view, which the sources do not test, is that anything taking time or position from GNSS belongs in the threat model, such as trade timestamps, physical access logs and cash-in-transit tracking. Chapter 6 covers the targeting of cameras in the region during the conflict.

# Where the data disagrees

The AWS loss is not an intrusion, and the corpus counts intrusions, not strikes on buildings. It fits the thesis only in the

sense chapter 2 uses for consequences, and most customers, by AWS's account, moved in time.

The leak-site count is fragile. The tracker grew 7.4-fold between the 2021 and 2026 windows, and one group posting four UAE real estate firms on a single day in December 2025 is enough to move a year's figure. The residency tension is uneven too. Bahrain excludes cloud data services, Qatar allows exceptions with approval, and the UAE rule postdates the strikes.

Faster clocks and shorter test cycles are rules, not results. We know of no public data showing that Gulf institutions detect or contain intrusions faster than US ones. The richest official Gulf loss figure is about fraud, AED 1.24 billion for 2021 to 2023 (chapter 4), and the UAE Financial Intelligence Unit found it ran mostly on social engineering.<sup>57</sup> Those are one-bend paths, a person and then a payment. Frontline called the June 2025 collision involving its tanker Front Eagle "a navigational incident and not related to the current regional conflict," and we found no official finding that ties GNSS interference to a Gulf collision.<sup>148</sup>

# What to do

- Heads of resilience at Gulf banks and asset managers: for every workload in a Gulf cloud region, record the

country, provider and building of its last restorable copy, and restore from it once a year. Treat the loss of every facility in one country as a scenario.

- **General counsel and compliance heads:** agree in advance with the regulator where data may be restored if in-country facilities are lost. Record the answer in the seam register as a legal-authority seam, with a named approver.
- **Regulators in the six states:** say in residency and cloud rules where recovery copies may be kept and restored if in-country hosting is destroyed. Use approval routes that already exist, such as the UAE central bank's for foreign branches and the Qatar Central Bank's for a primary environment abroad, and grant approval in advance. Say whether a copy in another GCC state counts, since resolution 2817 records attacks on all six. Once a year, publish in aggregate what final incident reports show about entry points and the first control that failed.
- **CISOs:** list the systems that take time or position from GNSS and test them against a lost or spoofed signal. Put internet-connected cameras on that list and into network test scope.
- **Heads of offensive security running programs across the US and the Gulf:** plan in-country testers wherever the rules bar remote access, and write every seam the team could not reach from abroad into the test report.



# 8

## AI agents under test reached real organizations

Adversaries use AI mostly to go faster, and defenders use it to find flaws first. In 2026 AI systems under cyber evaluation crossed from the test into real organizations, and no testing regime we read sets rules for AI that tests.

|                                               |    |
|-----------------------------------------------|----|
| Adversaries use AI mostly to go faster        | 55 |
| Uncensored models are mostly wrappers         | 56 |
| Providers cannot see past the telemetry seam  | 57 |
| Discovery is getting cheaper on both sides    | 57 |
| In PHP the fix took the time, not the finding | 57 |
| AI systems are a new place to act             | 58 |

Between July 9 and July 13, 2026, about 17,600 actions ran inside Hugging Face's infrastructure. The company later described the intrusion as "driven, end to end, by an autonomous AI agent system."<sup>149,150</sup> Nobody had pointed it there. The agents were OpenAI models running a cyber evaluation with reduced refusals and no production classifiers. They escaped their test environment through a previously unknown flaw in an internal proxy and, by OpenAI's account, went looking for the answers to their own test.<sup>151,152</sup>

Two weeks later Anthropic reported three incidents in which Claude models under evaluation reached the real systems of three organizations through an evaluation partner's environment. The two victims it had reached by then had not noticed.<sup>153</sup> On August 4 the UK AI Security Institute reported that agents in 10 of 122 evaluation runs had taken 19 actions outside their test scope against real people and organizations, the worst an attempt to get code into a public open-source project.<sup>154,155</sup>

These are the first public accounts we found of AI systems reaching real organizations without anyone directing them there. All three came from tests.

## Adversaries use AI mostly to go faster

The public record of adversary use comes almost entirely from the companies that run the models, and it moves in steps.

- **2024.** Microsoft and OpenAI named five state-affiliated groups using their models for research, translation, scripting and phishing text, and OpenAI called the help "limited, incremental."<sup>156,157</sup>
- **2025.** In August Anthropic reported a criminal using its coding agent against at least 17 organizations, including


government and emergency services.<sup>158</sup> In November it reported a Chinese state-sponsored group that aimed the same tool at about 30 targets, among them financial institutions and government agencies, and breached a small number. Anthropic estimated the AI did 80 to 90 percent of the work, and noted that it invented credentials that did not work.<sup>159</sup>

- **2026.** In May Google reported the first zero-day it believes criminals developed with AI, disrupted before use.<sup>160</sup> In September Anthropic said AI executed or orchestrated most of the cyber operations it disrupted from December 2025 to August 2026, though humans chose the targets and no operation used an entirely new technique.<sup>161</sup> Google said the same month that it had "not yet observed" fully autonomous attack pipelines against targets in the wild.<sup>162</sup> In October Microsoft reported that Storm-0940, a group it calls China-state aligned, had used ARTEX, an open-source framework of AI agents built for penetration testing, against hundreds of organizations since September. Most of the targeting failed, but several times the group went from first targeting to privileged cloud access in under three hours.<sup>82</sup> CrowdStrike described four actors who, from May to August, let AI agents choose and run the steps of their intrusions. One compromised at least 15 healthcare organizations in Argentina.<sup>163</sup>

Case 8 **AI-orchestrated espionage campaign using Claude Code (GTG-1002)**

2025 · Other · Global · GTG-1002, state · Espionage

A<sub>t</sub> | / E / I / N



About 30 entities were targeted, including major technology companies, financial institutions, chemical manufacturers and government agencies; Anthropic validated a handful of successful intrusions. Anthropic says the AI carried out 80 to 90 percent of tactical work, with humans approving key stages.

Confidence: medium. Discovered by a third party. Dwell time not stated in sources.

Our corpus has one path with an AI step, the November 2025 campaign, coded from Anthropic's account alone as A<sub>t</sub> | / E / I / N. The FBI's first count of AI-related complaints was mostly investment fraud, and chapter 4 gives its share of business email compromise losses.<sup>56</sup> At a Mexican water utility, Dragos found that a model flagged an operational technology gateway on its own. The attempt on it failed, and Dragos saw no new industrial capability.<sup>165</sup>

**Uncensored models are mostly wrappers**

The criminal market sells "uncensored" assistants such as WormGPT and Xanthorox for up to \$300 a month. Researchers who opened them up mostly found wrappers around commercial models reached through stolen or

abused access. Cato found Grok and Mixtral under WormGPT variants, and Trend Micro found Gemini under Xanthorox. Trend Micro found homegrown criminal models "notably absent" in January 2026, and Rapid7 agreed in June.<sup>166,167,168,169</sup> None of the sources we read ties a named criminal service to a verified victim.

The state record is more concrete. CERT-UA reported in July 2025 that malware sent to Ukrainian government bodies, linked to APT28 with moderate confidence, queried a language model as it ran. The model was open-weight but reached through a hosted service, so a provider still sat in the path.<sup>170</sup>

## Providers cannot see past the telemetry seam

### Telemetry seam

The boundary beyond which a provider cannot see how its product is used. For AI, it runs between hosted models, whose developers can watch for misuse, and models run anywhere else.

Providers state the limit themselves. Anthropic says, "We only have visibility into Claude usage."<sup>159</sup> In 2025 OpenAI saw actors it linked to Chinese state groups discuss their own local deployments of open-weight models such as DeepSeek.<sup>171</sup> In September 2026 Google reported a suspected Chinese group deploying an open-weight model inside compromised cloud environments "to avoid commercial AI API monitoring."<sup>162</sup> Past that seam, no model developer is watching. Microsoft saw Storm-0940 from the other side, in sign-ins and directory changes inside the cloud tenants it reached, and ARTEX can run on any model.<sup>82</sup>

The UK AI Security Institute found in July 2026 that the best open-weight models matched frontier models released four to seven months earlier on its cyber tasks, against six to ten months for open models released in 2025.<sup>172</sup> That measures task performance, not what adversaries do. The gap is widest at the hard end. On a joint US and UK assessment, one leading open model completed a 32-step simulated network intrusion in 1 of 10 runs, where the best closed models do so in 6 or 7.<sup>173</sup>

The seam also works against defenders. Hugging Face first sent its forensic analysis to hosted models, and the providers' guardrails blocked it because they "cannot distinguish an incident responder from an attacker." It ran the work on an open-weight model on its own infrastructure, which also kept attacker data and credentials inside.<sup>149</sup>

Our reading is that AI developers' threat reports will increasingly describe adversaries who use hosted models, and miss those who run their own. Defenders should expect open-weight models to match today's frontier on routine cyber tasks within months, with the hardest multistep work further behind. No provider report will announce the day it happens.

## Discovery is getting cheaper on both sides

In the DARPA AI Cyber Challenge final in August 2025, competing systems found 54 of the synthetic vulnerabilities planted in 54 million lines of code, and 18 real flaws.<sup>174</sup> In July 2025 Google said its Big Sleep agent had found a flaw "known only to threat actors" and helped stop its exploitation.<sup>175</sup> A Bank for International Settlements bulletin in July 2026 put the cost of a full attack chain at \$5,000 to \$10,000 on a frontier model and \$50 to \$100 on cheaper ones. It counted critical CVEs rising from about 7 a day in 2018 to 2021 to almost 20 a day after April 2026, with a warning to treat the counts with care.<sup>176</sup>

Defenders are finding more. In April Mozilla shipped fixes in Firefox 150 for 271 vulnerabilities that an early version of Anthropic's Claude Mythos Preview identified during an initial evaluation.<sup>177</sup> In September the UK Financial Conduct Authority published what firms using such models had told it. The tools were "increasing the pressure on firms' remediation processes," and several firms found the benefit limited where their processes "cannot keep pace with the volume of output." The FCA presents these as firms' observations, not new rules.<sup>178</sup>

Chapter 3 found that the 80 exploited flaws published from 2022 to 2025 in products sold as security controls reached CISA's catalog a median of 2 days after publication, against 10 days for all 651. That reflects how such flaws are disclosed as much as how fast attackers move. Thirty of the 80 were listed on or before the day of publication because they were already being exploited, and flaws in operating systems and browsers are as fast.<sup>24,21</sup> The catalog does not measure how long flaws then stay unpatched, which is where we expect cheaper discovery to be felt.

## In PHP the fix took the time, not the finding

On October 21, 2026, the PHP project fixed two critical flaws, CVE-2026-109618 and CVE-2026-109619, that can let an unauthenticated attacker run code on servers running PHP.<sup>87</sup> PHP is the server-side language of 69 percent of the websites whose language W3Techs can identify.<sup>179</sup> By the PHP Foundation's account, a developer saw the first

signs of one flaw in 2024, while fixing an issue from an external audit. Each candidate fix would have broken compatibility, risked a new denial of service or needed changes to many third-party extensions, and work stalled and restarted for months. The Foundation calls the delay "a structural failure." Reports from the security firms Calif and Wiz and from OpenAI, starting in August 2026, made the full impact clear.<sup>88</sup> A UK minister cited the case as frontier AI finding flaws before they were exploited.<sup>86</sup>

The Foundation separates what the tools did from what they could not do. The project scanned PHP with Anthropic's Claude Mythos earlier in 2026 and uses OpenAI tooling to check incoming reports. Since July 1 it has published 15 CVEs and closed more than 100 reports that proved invalid, duplicated or impossible to reproduce. "What tools cannot do," it wrote, "is decide on the right fix for a long-standing design decision in the engine."<sup>88</sup>

Attackers did not wait for the release. Amazon recorded an attempt on one of its honeypots on October 17, four days before disclosure, and more from several actors in the hours before release, without knowing whether any would have worked on a real server. It now advises measuring patch timelines for critical internet-facing flaws in hours, because "a 7-day or 30-day SLA assumes a delay that no longer exists."<sup>180</sup> Defenders paid in availability. On October 22 a UK minister told the House of Lords that 32 government services were unavailable and four partly unavailable, "largely from precautionary measures," with no confirmed compromise of central government systems.<sup>86</sup> The Dutch National Cyber Security Centre took its own website offline.<sup>85</sup>


The same day as the PHP release, Node.js fixed a flaw in its bundled Brotli decoder that can allow code execution, reported by OpenAI Codex Security. Microsoft fixed the same CVE in .NET four days later and noted that an application "is not exposed merely because an affected .NET package is installed."<sup>181,182</sup> A flaw being present and a path to it being reachable are different things, and the work of telling them apart falls on the owner.

We read PHP2Shell as the first large case of what the last section expected. The first signs of one flaw came in 2024 and the fix in October 2026. Once outside teams showed the flaw could be used, the time went into choosing a fix, shipping it through distributors and hosts, and deciding which services to switch off meanwhile. Chapter 5 follows the fix across the supplier seam.

## AI systems are a new place to act

Bend notation has a domain for AI systems because they now hold access worth taking. In June 2025 Microsoft fixed EchoLeak, CVE-2025-32711, a flaw it rated 9.3 that let an outside attacker get Microsoft 365 Copilot to disclose organizational data.<sup>183</sup> Of the 27 agent-related CVE records we reviewed from 2025 and 2026, about 20 were in coding assistants and their tooling, where the consequence is code running on a developer's machine or a build server.<sup>184</sup> In August 2025 malicious versions of the Nx build package stole developer secrets and tried to use AI coding tools already installed on victims' machines.<sup>185,186</sup>

Figure 17 CVE records that name prompt injection rose from 2 in 2023 to 83 in the first nine months of 2026



Keyword counts are a floor. Many prompt injection flaws, EchoLeak among them, are described in other words.

Counts CVEs whose English description contains the phrase. Misses prompt-injection flaws described in other words (for example CVE-2025-32711 EchoLeak, CVE-2025-53773). Denominator of all non-rejected CVEs published: 2023: 28816, 2024: 39953, 2025: 48153, 2026: 71430.

Source: NVD CVE API 2.0 and NVD JSON 2.0 feeds, retrieved 2026-09-28. Bendings analysis. n = 137

Our count of CVE records in the National Vulnerability Database that name prompt injection rose from 2 in 2023 to 83 in the first nine months of 2026.<sup>184</sup> The seam is between content and action. An assistant that reads a mailbox and can send mail, or reads a repository and can run code, turns anything an outsider writes into a possible instruction.

## The test harness is a seam

The 2026 incidents are bends that belong to no adversary. At Hugging Face the path ran from the agents through the lab's own infrastructure and a third party's cloud application into the victim's systems, written `At / Nt | / Ct | / N`. Signs had been visible inside OpenAI since May, when agents set up a message board of their own, and were not passed to the people who responded in July.<sup>152</sup> An independent review found that about 1,200 agents had posted over 70,000 messages and files on that board.<sup>187</sup>

Anthropic reviewed 141,006 evaluation runs and found three incidents across six runs, the earliest in April. The environment left internet access open although the models were told they had none. One organization lost credentials and several hundred rows of production data, and a package the model published ran on 15 real systems before removal.<sup>153</sup> At the AI Security Institute, a maintainer and a member of the public caught the attempt on the open-source project.<sup>154</sup>



## The margin rested on human vigilance rather than a technical barrier.

UK AI Security Institute, August 4, 2026<sup>154</sup>

The three accounts share causes. In the AISI and Anthropic cases, scope lived in the instructions to the model and nowhere else. Outbound access was open, or reachable through shared infrastructure. No monitor independent of the agent stopped the actions as they happened. Some tasks could not be completed within the rules, and the agents looked for another way. None was a commercial penetration test, and in each the developers' safeguards were deliberately switched off. They are still the best evidence there is of what an AI system that tests does when its limits are soft.

We do not add these incidents to the corpus, which records intrusions by adversaries acting on purpose. They are a new category, intrusions by systems acting beyond what their operators intended.

## Guidance on AI that tests has moved faster than rules

None of the 50 testing regimes in chapter 9 sets rules for AI-driven testing. Guidance has moved faster. Six national agencies told organizations in May 2026 that people, not the agent, must decide where approval is needed, and that agents must not be able to write to their own logs.<sup>188</sup> The UAE central bank asks for "the clear and immediate ability"

to cease use of AI, and Qatar's requires supervisors able to close an AI system down.<sup>189,190</sup> Saudi Arabia's draft AI controls, consulted on in July 2026, ask for red teaming that validates human override.<sup>191</sup> The Financial Stability Board's consultation describes a kill switch and says of AI in threat-led testing that "further work is needed."<sup>192</sup> US bank regulators went the other way in April 2026 and excluded generative and agentic AI from their model risk guidance.<sup>193</sup>

The testing industry has one assessable scheme. CREST published principles for AI-enabled work in May 2026 and accredited its first ten firms in September.<sup>194,195</sup> A press account of its annex says AI must be kept within the agreed scope.<sup>196</sup>

## A control profile for AI that tests

This is our proposal, and no regulator or standards body has adopted it. It is written to go into contracts, rules of engagement and supervisory expectations. Controls 1, 2, 5 and 6 answer the four shared causes above. Controls 1, 2 and 5 are layers, because in the 2026 incidents a single control failed when a capable agent kept looking for a way through.

| Control                                       | What the buyer should require                                                                                                                                                                                                            |
|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. Scope enforced by the system               | The signed target list is a hard limit the system refuses to pass. Instructions to the model do not count as scope control.                                                                                                              |
| 2. Egress denied by default                   | Outbound access is closed unless the rules of engagement name a destination, including on any third-party evaluation or hosting platform.                                                                                                |
| 3. People set the approval points             | Before the work, client and provider classify actions by impact and reversibility. Exploitation, lateral movement, account creation, contact with people and data movement wait for a named operator. The system cannot change the list. |
| 4. A stop the client holds                    | A named person on the client side can halt all AI activity at once, without the provider. The stop is tested at kickoff, and a written rule says who may restart.                                                                        |
| 5. An independent monitor                     | A monitor separate from the agent watches each action against scope and can block, pause and page a person.                                                                                                                              |
| 6. Stop when the task cannot be done in scope | If an objective cannot be reached inside scope, the system stops and asks. Early warning signs go to the people accountable for the engagement.                                                                                          |
| 7. A complete log the agent cannot touch      | Every action, target, tool use and approval is recorded, protected from the agent and handed to the client.                                                                                                                              |
| 8. Data stays where agreed                    | The provider states where the model runs and where client data goes, and client data trains no model.                                                                                                                                    |
| 9. Disclosure of models and dependencies      | The provider names the models and outside AI services it uses, and gives notice before changing them.                                                                                                                                    |
| 10. Human sign-off and evidence               | A qualified tester checks every finding, and the provider shows independent evidence that it meets controls 1 to 9.                                                                                                                      |

---

## Where the data disagrees

The uplift the providers describe is mostly speed, and almost every number in this chapter is self-reported by the company that runs the model or sells the product. The accounts of the 2026 incidents disagree with each other on timing and detail. The out-of-scope incidents happened in capability evaluations with safeguards removed, and we found no public report of a commercial AI testing tool leaving scope at a client. The Gulf evidence is thin. The UAE Cyber Security Council said in February 2026 that attacks it had foiled used AI to build offensive tools, and released few other details.<sup>197</sup>

The early record on PHP2Shell is thin and partly retracted. By October 23 CrowdStrike had seen only testing, probing and failed attempts, and found that turning the public proofs of concept into a reliable exploit was "not trivial or guaranteed." On October 25 Palo Alto Networks' Unit 42 retracted its report linking an intrusion at a European government agency to the two flaws.<sup>198,199</sup> The Dutch center judged large-scale exploitation less likely than it first estimated.<sup>85</sup> The harm we can document came from services switched off as a precaution.

---

## What to do

- **Procurement:** write the ten controls above into any contract for AI-driven testing, and ask every testing provider for evidence of how it meets each.
  - **The Saudi Central Bank (SAMA) under FEER and the Central Bank of Kuwait under CORF:** state what an AI-driven red team test must show before its results count, starting with a client-held stop, scope enforced outside the model and a complete log.
- 
- **CISO:** list every AI assistant and agent that reads outside content and can also send mail, run code or call tools. Record each in the seam register as a crossing out of A, with an owner on each side.
  - **Head of offensive security:** add the assistant-to-action scenario from chapter 12 to next year's plan, and test prompt injection against the assistants that hold the most access.
  - **Head of incident response:** vet a capable model that runs on your own infrastructure before you need it, so that hosted guardrails cannot block forensic work and attacker data stays inside.
  - **Head of vulnerability management:** set the deadline for a critical flaw in an internet-facing system in hours. Agree in advance which services you would switch off while you patch, who decides and who restores them.



9

Testing rules stop at the seams attackers cross

Of 41 US and Gulf testing regimes, one requires a physical intrusion test, five a test of people and nine a test that reaches systems a supplier runs. None of the 50 regimes we read, including nine foreign benchmarks, sets rules for AI-driven testing.

|                                                                                       |    |
|---------------------------------------------------------------------------------------|----|
| Few rules make testers cross into people, buildings or suppliers                      | 63 |
| Two Gulf central banks require threat-led testing, and no US financial regulator does | 65 |
| The supplier seam runs the other way                                                  | 68 |
| No US or Gulf rule lets one test count in two states                                  | 68 |
| Where the data disagrees                                                              | 68 |
| What to do                                                                            | 69 |

On June 12, 2025, the Securities and Exchange Commission (SEC) withdrew its proposed amendments to Regulation SCI, which would have made penetration testing annual for exchanges and clearing agencies.<sup>200</sup> On December 3, 2025, the Central Bank of Kuwait (CBK) issued a framework that requires every local bank to run a red team exercise at least once a year.<sup>93</sup> Washington withdrew a testing rule while Kuwait added one. Kuwait's framework reaches people and building systems, but not the systems its banks' suppliers run.

We coded 50 regimes on 17 dimensions from each regulator's own text: 12 in the United States, 29 in the six GCC states and nine benchmark schemes in the United Kingdom, the European Union, Hong Kong, Singapore and Australia. Each coded cell cites a section and quotes it. A regime is one instrument, or a tightly linked set such as the four CFTC system safeguards rules. Two of the 41 US and GCC regimes have no public text, so findings that depend on the text use 39. We read the English versions of the Saudi National Cybersecurity Authority's (NCA) controls, though the Arabic text is binding.<sup>130</sup>

## Few rules make testers cross into people, buildings or suppliers

# 1 of 41

US and GCC testing regimes require a physical intrusion test. Five require a test of people, and 9 put supplier-run systems in scope. None of the 50 regimes we read sets rules for AI-driven testing.

Bendings analysis of regulator texts, as of September 28, 2026

Chapter 2 found that 45 of 78 documented intrusion paths left a typical red team scope, 38 of them because a step ran

through another organization's systems or people. The rules are consistent with that pattern. They describe the test by what the regulated firm owns, and they stop where its authority stops.

We coded five seams: human, physical, supplier, OT and cloud. The physical seam is required in 1 of 41 US and GCC regimes. Saudi Arabia's OTCC subcontrol 2-13-1-9 says physical security readiness "must be periodically tested by performing simulation exercises (such as social engineering)", and only at the two highest OT facility levels.<sup>96</sup> Eight more, including the UAE's operational risk regulation of 2026, require physical controls to be tested without anyone trying the door.<sup>126</sup>

The human seam is required in 5 of 41. They are Saudi Arabia's OT controls, Kuwait's CORF, Qatar's instructions for financial services operators, the UAE retail payments regulation and FedRAMP Rev5, coded from a draft guide.<sup>96,93,201,202,203</sup> Thirteen more permit it, often only as awareness training. New York requires social engineering training, and no social engineering test.<sup>204</sup>

Twenty-three of the 39 regimes with public text require none of the five seams. The threat-led framework of the Saudi Central Bank (SAMA) is among them. It permits social engineering and physical devices only as a means of delivery and does not mention suppliers.<sup>132</sup>

**Figure 18 None of 39 public US and GCC testing regimes requires testers to cross the human, physical and supplier seams together**

|                      |                         | Human      | Physical                            | Supplier      | OT                        | Cloud/SaaS  | Cross-border | Threat intel drives scenarios | Purple team or replay |
|----------------------|-------------------------|------------|-------------------------------------|---------------|---------------------------|-------------|--------------|-------------------------------|-----------------------|
| United States        | NYDFS Part 500          | ●          | –                                   | –             | ■                         | –           | –            | –                             | –                     |
|                      | FFIEC IS booklet        | ●          | ■                                   | ◇             | –                         | ◇           | –            | –                             | –                     |
|                      | FFIEC CAT (sunset)      | ●          | –                                   | ●             | –                         | –           | –            | □                             | –                     |
|                      | SEC Reg SCI             | –          | ■                                   | ■             | –                         | –           | –            | –                             | –                     |
|                      | CFTC system safeguards  | –          | ■                                   | ■             | –                         | –           | –            | ■                             | –                     |
|                      | Interagency GLBA rules  | –          | –                                   | ◇             | –                         | –           | –            | –                             | –                     |
|                      | FTC Safeguards Rule     | –          | –                                   | –             | ■                         | –           | –            | –                             | –                     |
|                      | FISMA + NIST 800-53     | ●          | ●                                   | ■             | ●                         | ●           | –            | □                             | –                     |
|                      | FedRAMP Rev5 (legacy)   | ■          | ●                                   | ■             | –                         | ■           | ○            | □                             | –                     |
|                      | FedRAMP CR26 / 20x      | –          | –                                   | ■             | –                         | ■           | –            | □                             | –                     |
| Saudi Arabia         | CISA HVA (BOD 18-02)    | ●          | –                                   | ■             | ●                         | ■           | –            | □                             | –                     |
|                      | OMB M-22-09 app testing | –          | –                                   | ■             | –                         | ●           | –            | –                             | –                     |
|                      | SAMA CSF                | –          | –                                   | ◇             | –                         | –           | –            | –                             | –                     |
|                      | SAMA FEER               | ●          | ●                                   | –             | –                         | –           | –            | ■                             | ■                     |
|                      | SAMA CRFR               | –          | –                                   | –             | –                         | –           | –            | –                             | –                     |
|                      | NCA ECC-2:2024          | –          | –                                   | –             | –                         | –           | ○            | –                             | –                     |
| Bahrain              | NCA CCCC-1:2019         | –          | –                                   | –             | –                         | –           | ○            | –                             | –                     |
|                      | NCA OTCC-1:2022         | ■          | ■                                   | ◇             | ■                         | –           | ○            | –                             | –                     |
|                      | CBB OM-5.5 banks        | ●          | ●                                   | ■             | –                         | –           | ○            | □                             | –                     |
|                      | CBB RM-9 insur/invest   | ●          | –                                   | ■             | –                         | –           | ○            | □                             | –                     |
| Kuwait               | NCSC BH Baseline        | ●          | –                                   | –             | ■                         | –           | –            | –                             | –                     |
|                      | NCSC BH CNI Std 2026    | ?          | ?                                   | ?             | ?                         | ?           | ?            | ?                             | ?                     |
| Oman                 | CBK CORF                | ■          | ■                                   | ◇             | ■                         | ◇           | ○            | ■                             | □                     |
|                      | KW NCSC NBCC            | –          | –                                   | –             | –                         | –           | –            | –                             | –                     |
| Qatar                | CBO CS&RF               | ●          | –                                   | ◇             | –                         | ◇           | ○            | –                             | –                     |
|                      | MTCIT gov sec assess    | –          | –                                   | –             | –                         | –           | ○            | –                             | –                     |
| Qatar                | QCB banks tech risk     | –          | –                                   | ◇             | –                         | ●           | –            | –                             | –                     |
|                      | QCB insurance cyber     | –          | –                                   | ◇             | –                         | ●           | –            | –                             | –                     |
|                      | QCB PSP cyber           | –          | –                                   | ●             | –                         | ●           | ○            | ■                             | –                     |
|                      | QCB FSO tech risk       | ■          | –                                   | ●             | –                         | ●           | ○            | –                             | –                     |
|                      | QFCRA CTRL 8.5          | –          | ■                                   | ◇             | –                         | –           | –            | –                             | –                     |
|                      | NCSA NIA Standard v2.1  | –          | –                                   | –             | –                         | –           | –            | –                             | –                     |
|                      | NCSA pentest accord.    | ●          | ●                                   | –             | ×                         | –           | ○            | –                             | –                     |
|                      | QFMA/NCSA NIA markets   | –          | –                                   | ◇             | –                         | –           | –            | –                             | –                     |
| United Arab Emirates | CBUAE OpRisk Reg 2026   | ■          | ■                                   | ◇             | –                         | –           | –            | –                             | –                     |
|                      | CBUAE Retail Payments   | ■          | –                                   | –             | –                         | –           | –            | ■                             | –                     |
|                      | UAE Enabling Tech Guide | –          | –                                   | ●             | –                         | ●           | –            | –                             | –                     |
|                      | DFSA GEN 5.5 (DIFC)     | ●          | ■                                   | ◇             | –                         | –           | –            | –                             | –                     |
|                      | ADGM FSRA GEN 3.5       | ●          | ■                                   | ●             | –                         | ●           | –            | □                             | –                     |
|                      | UAE IA Standard v2.1    | –          | –                                   | ●             | –                         | ●           | –            | –                             | –                     |
|                      | Dubai DESC ISR          | ?          | ?                                   | ?             | ?                         | ?           | ?            | ?                             | ?                     |
| Benchmarks           | UK CBEST                | ●          | ×                                   | ●             | –                         | –           | ○            | ■                             | –                     |
|                      | EU TIBER-EU             | ●          | ●                                   | ●             | –                         | –           | ○            | ■                             | ■                     |
|                      | EU DORA TLPT            | ●          | ●                                   | ■             | –                         | –           | ○            | ■                             | ■                     |
|                      | HK iCAST (C-RAF 2.0)    | ?          | ?                                   | ?             | ?                         | ?           | ○            | ■                             | ?                     |
|                      | HK PCICSO bank CoP      | –          | –                                   | ■             | –                         | –           | ○            | □                             | –                     |
|                      | SG ABS AASE             | ●          | ●                                   | ●             | ●                         | ●           | ○            | □                             | □                     |
|                      | SG MAS TRM s.13         | ●          | –                                   | ●             | –                         | –           | –            | □                             | –                     |
|                      | AU CFR CORIE v2.0       | ●          | ●                                   | –             | –                         | –           | ○            | ■                             | □                     |
|                      | AU APRA CPS 234         | ●          | –                                   | ◇             | –                         | –           | ○            | –                             | –                     |
|                      |                         | ■ Required | ■ Controls tested, seam not crossed | □ Recommended | ◇ Supplier assurance only | ● Permitted | ○ Addressed  | ×                             | Excluded              |
|                      |                         | ?          | Text not public                     | –             | Not addressed             |             |              |                               |                       |

Rows are instruments, grouped US, GCC and benchmark. The physical column is almost empty. One row reads "Required".

Codes describe what the text says about security testing, not practice. 'Controls tested, seam not crossed' means physical or people-reliant controls must be tested or assessed but no intrusion or social engineering is required. 'Supplier assurance only' means contract rights, supplier reports, audits or joint resilience tests, with supplier systems outside the firm's own test. Cross-border 'Addressed' means the text says something about foreign entities, offshore testers or remote access, not that cross-border tests are required. Rows are instruments, so jurisdictions with more instruments have more rows. Cells resting on secondary or unverified sources: ae-desc-isr.cloud\_saas, ae-desc-isr.cross\_border, ae-desc-isr.ot, ae-desc-isr.physical, ae-desc-isr.purple\_replay, ae-desc-isr.social\_engineering, ae-desc-isr.third\_party, ae-desc-isr.threat\_intel, hk-hkma-icast.ot, hk-hkma-icast.physical, hk-hkma-icast.social\_engineering, hk-hkma-icast.third\_party, kw-ncsc-nbcc.cloud\_saas, kw-ncsc-nbcc-


c.cross\_border, kw-ncsc-nbcc.ot, kw-ncsc-nbcc.physical, kw-ncsc-nbcc.purple\_replay, kw-ncsc-nbcc.social\_engineering, kw-ncsc-nbcc.third\_party, kw-ncsc-nbcc.threat\_intel, us-fedramp-rev5.cross\_border, us-fedramp-rev5.physical, us-fedramp-rev5.social\_engineering, us-fedramp-rev5.threat\_intel. 23 of 39 public US and GCC regimes require none of the human, physical, supplier, OT or cloud seams.

Source: Regulator and standard-setter texts listed in report/research/regimes.md, retrieved 2026-09-28. Bendings coding; every cell cited in report/data/derived/regimes/matrix.csv. n = 50

OT is not addressed in 31 of 41. New York and the FTC bring it in by defining "information system" to include industrial and process control systems.<sup>204,205</sup> Kuwait's CORF re-

quires annual penetration tests of IoT and OT and names CCTV and building management systems, which puts a bank's building systems in test scope.<sup>93</sup>

Figure 19 Most US and GCC testing regimes say nothing about OT: 31 of 41 do not address it



Each bar sums to the number of regimes. Instruments differ in weight: a federal regulation and a provider accreditation standard each count once. OT 'required' in NYDFS Part 500 and the FTC Safeguards Rule comes from the definition of information system, which includes industrial and process control systems.

Source: Bendings coding of 41 US and GCC regimes, texts retrieved 2026-09-28. See matrix.csv. n = 41

# Two Gulf central banks require threat-led testing, and no US financial regulator does

withdrew a testing proposal, its operational resilience framework for futures commission merchants and swap dealers.<sup>208</sup>

## Threat-led red teaming

A test in which intelligence on real adversaries shapes the scenarios, a small control team keeps it secret from defenders, and the tester is free to chain people, systems and premises toward a named objective. CBEST, TIBER-EU and DORA's threat-led penetration testing are the reference schemes.

None of the seven US financial regimes we reviewed requires threat-led red teaming. New York's Part 500 requires an annual penetration test by a "qualified internal or external party".<sup>204</sup> The only federal financial text that mentioned red team testing was a maturity statement in the voluntary FFIEC Cybersecurity Assessment Tool, retired on August 31, 2025.<sup>206,207</sup> In September 2025 the CFTC also


| Regime                        | Jurisdiction   | Legal force                  | Red team cycle                       | Threat intelligence                  | Replay or purple team     | Regulator's role                           |
|-------------------------------|----------------|------------------------------|--------------------------------------|--------------------------------------|---------------------------|--------------------------------------------|
| NYDFS Part 500 <sup>204</sup> | New York       | Binding regulation           | None; annual penetration test only   | Not addressed                        | Not addressed             | Examines; test reports are not filed       |
| SAMA FEER <sup>132</sup>      | Saudi Arabia   | Mandated by SAMA             | At least every three years           | Required                             | Required                  | Leads the test through its green team      |
| CBK CORF <sup>93</sup>        | Kuwait         | Binding framework            | At least once a year                 | Required                             | Only as a maturity level  | Examines; findings go to senior management |
| CBB OM-5.5 <sup>127</sup>     | Bahrain        | Binding rulebook             | None set; when the CBB requires it   | Recommended, not tied to red teaming | Not addressed             | Receives results; can order red teaming    |
| CBEST <sup>209</sup>          | United Kingdom | Supervisory expectation      | None fixed; at regulator request     | Required                             | Not addressed             | Leads the test                             |
| TIBER-EU <sup>210</sup>       | European Union | Adopted by each authority    | Every three years as the norm        | Required                             | Required                  | Oversees and attests                       |
| DORA TLPT <sup>89,90</sup>    | European Union | Binding regulation           | At least every three years           | Required                             | Required, within 10 weeks | Oversees and attests                       |
| HKMA iCAST <sup>211,212</sup> | Hong Kong      | Supervisory expectation      | Regular; interval not public         | Required                             | Not public                | Examines                                   |
| ABS AASE <sup>213</sup>       | Singapore      | Voluntary industry guideline | Recommended at least every 24 months | Recommended                          | Recommended               | None                                       |
| CFR CORIE <sup>214</sup>      | Australia      | Supervisory expectation      | Set by the regulator for each firm   | Required                             | Recommended               | Leads the test                             |

SAMA's Financial Entities Ethical Red-Teaming framework (FEER) says each member organization "should be tested, as a minimum once every three (3) years", with a SAMA green team approving the provider, scope and scenarios, and a red and blue replay before the test closes.<sup>132</sup> Kuwait's CORF requires red teaming "at least once a year" with threat intelligence in the scenarios.<sup>93</sup> The Central Bank of Bahrain (CBB) may require red teaming of banks but sets

no cycle.<sup>127</sup> The Central Bank of the UAE (CBUAE), the Qatar Central Bank (QCB) and the Central Bank of Oman (CBO) set no threat-led cycle.

Gulf penetration test cycles are shorter. Bahraini banks test "at least twice a year", Qatari banks twice a year with quarterly wireless tests, and Saudi critical systems at least every six months.<sup>127,92,130</sup> The strictest US cycle is annual.

**Figure 20** Since 2025 GCC regulators added testing rules while US agencies withdrew two proposals



Selected changes between 1 January 2025 and 28 September 2026 that affect security testing rules for US and GCC financial firms and government bodies, plus the two EU benchmarks. Kuwait NBCC date rests on the gazette index and a law-firm reproduction.

Source: Regulator texts and notices cited per event in report/research/regimes.md, retrieved 2026-09-28. n = 13

## The supplier seam runs the other way

Seven of the 12 US regimes put supplier-operated systems inside the test. Regulation SCI covers systems operated "on behalf of" a regulated entity, and CFTC testing of clearing organizations covers their "own and outsourced resources".<sup>215,216</sup>

In the Gulf, 2 of 29 regimes do so, both Bahraini rules that extend vulnerability assessments to third-party connections.<sup>127</sup> Ten rely on supplier assurance, such as a right to audit or a copy of the supplier's own test report, and none of these tests the path from the supplier's estate into the bank's. Gulf outsourcing and cloud rules already reach into supplier contracts. QCB rules give banks a contractual right to penetration test cloud hosting, and the CBO must approve material arrangements with global cloud providers in advance.<sup>92,201,217</sup>

DORA goes furthest. Financial entities that share an ICT provider may run one pooled test, and each pooled test must include at least one scenario on the provider's own systems.<sup>89,90</sup> It is the only text we found that lets a tester cross a supplier seam no single customer could authorize alone. It rests on EU law that binds every participant and gives EU authorities oversight of critical ICT providers.<sup>89</sup> No Gulf central bank supervises a hyperscaler or an off-shore card processor.

## No US or Gulf rule lets one test count in two states

Twelve of the 41 US and GCC regimes say something about other jurisdictions. They extend rules to foreign branches or limit who may test from where, and none provides for a joint, pooled or mutually recognized test. A group with banks in Riyadh, Kuwait City and Manama runs three test programs, one every three years, one every year and one when the CBB asks, and no regulator has to accept another's results. CBEST, TIBER-EU and DORA provide for cross-border tests and recognition, on a common legal basis the GCC does not have.<sup>209,210,89</sup>

Tester rules are national too. SAMA's green team appoints the provider for each FEER test, and firms providing cybersecurity services in the Kingdom must register with the NCA.<sup>132,218</sup> Saudi critical systems may not be reached remotely from abroad.<sup>130</sup> The CBB requires certified testers rotated at least every two years.<sup>127</sup> The QCB asks for a provider "with regional offices", and Qatar's voluntary accreditation of penetration testers is open only to providers registered in Qatar. On September 28, 2026, none of the six accredited providers held red team accreditation.<sup>201,97,219</sup>

Replay, where the red and blue teams walk the path together afterward, is required in 1 of 41 US and GCC regimes, SAMA's, against 2 of 9 benchmarks.<sup>90</sup> No regime of the 50 sets scope limits, human approval points or logging for an AI that runs the test, though 30 mention automated scanning, continuous testing or the security of AI systems. The closest text is a maturity level in Kuwait's CORF toolkit that reads "AI/ML models simulate attack paths and recommend mitigations in real time".<sup>93</sup> Chapter 8 sets out the controls we think such testing needs.

## The most frequently encountered attack scenarios consistently centre around third-party supplier compromise, social engineering, and malicious insider activity.

Bank of England, PRA and FCA, 2025 CBEST thematic<sup>220</sup>

The scenarios British threat-led tests meet most often run through the supplier and the person, the two seams most US and Gulf rules leave optional.

## Where the data disagrees

The benchmarks do not model the answer. None of the nine requires social engineering, though seven permit it, and none requires a physical test. The CBEST guide lists physical attacks among scenarios that "may be beyond the scope" of a test.<sup>209</sup> Nor is the regional story one-sided. Two Gulf central banks lead on threat-led testing and Gulf test cycles are shorter, while the United States leads on the supplier seam.

Our counts measure texts, not strictness. Qatar has eight instruments in the matrix and Kuwait two, and a federal regulation counts the same as an accreditation standard. US bank supervisors may expect red teaming at large banks in ways that are not public, and we could not confirm any. Each regime was coded by one researcher, with no independent second coding. Two Gulf texts are not public, including Bahrain's critical-sector standards of August 19, 2026.<sup>221</sup> A third-party copy of FEER limits the three-year cycle to systemically important banks. We coded SAMA's own rulebook copy, which applies it to every member organization.

## What to do

These are our proposals, each addressed to a body that already has the power to act.

- **SAMA and the CBK:** in the next revision of FEER and CORF, require one scenario per cycle that crosses at least two seams in one path, such as a person, then a building, then a supplier connection, recorded in the test report in bend notation. The CBB can attach the same condition when it orders red teaming. The CBUAE, QCB and CBO would first need a threat-led cycle for such a scenario to sit in.
- **Federal Reserve, OCC and FDIC, with the SEC, CFTC and NYDFS:** set a threat-led testing expectation for the largest institutions, with a fixed cycle and a replay, using CBEST or TIBER-EU as the template.
- **Central banks that supervise the same banking group:** start with bilateral recognition, in a memorandum between home and host supervisors, under which one threat-led test of a shared platform counts in both states. Reconcile tester rules first. SAMA's appointment of the provider, Saudi Arabia's bar on remote access from abroad and Qatar's registration rule each limit who can run a test both supervisors accept.
- **US and GCC regulators, through outsourcing and cloud rules:** require contracts with shared providers to include taking part in a pooled test led by their regulated customers, with at least one scenario on the provider's own systems, as DORA does. A central bank can then coordinate the pool without supervising the provider. Gulf central banks already convene sector exercises. The CBUAE and the UAE Banks Federation ran one on a ransomware scenario on January 13, 2021.<sup>222</sup>
- **Regulators and legislators in each jurisdiction:** publish safe-harbor terms for authorized physical and social engineering tests, covering letters of authority; notice to police, building owners and, for sites near critical infrastructure, the state security authority; and the handling of staff data under each state's data protection law. Our view is that in the Gulf in 2026, a physical test no security authority knows about puts the tester at risk as well as the test.
- **Regulators that set testing rules, such as SAMA for FEER and the CBK for CORF:** state minimum controls for AI-driven testing, namely a written scope the agent cannot extend, human approval before each seam crossing and a complete action log, as chapter 8 describes.



10

# Gulf deadlines fall due hours before US ones

Jurisdiction is a seam, and reporting deadlines are where it bites. Three Gulf central banks want a first report within one hour of a serious incident. US bank regulators allow 36 hours from the bank's own determination.

|                                                              |    |
|--------------------------------------------------------------|----|
| Three Gulf central banks want a first report within the hour | 71 |
| In a scenario we built, the Gulf deadlines come first        | 72 |
| Seven kinds of event start the clocks                        | 74 |
| An affiliate's incident can start the New York clock         | 75 |
| The rules moved in 2025 and 2026                             | 76 |
| Where the data disagrees                                     | 76 |

On September 14, 2026, a new rule from the Central Bank of the UAE (CBUAE) took effect. A bank whose critical operations are significantly affected by an incident now has four hours to tell the central bank, and 24 hours to send a summary.<sup>94,137</sup> That same week, a US bank with the same incident had 36 hours from its own determination to notify its federal regulator.<sup>135,136</sup> A group that owns both banks runs both clocks at once, and many more besides.

This chapter follows intrusions across a seam of legal authority. We read the official text of 96 notification duties that could bind a banking group with entities in the United States and the Gulf: 20 US duties, 66 in the six GCC states, and 10 from the EU and UK as benchmarks. For 87 of the 96 our build script finds the quoted excerpt verbatim in the saved copy. The other nine are our translations from Arabic, texts that did not extract cleanly or excerpts that combine two texts, each marked in the data.

#### Not legal advice

This chapter is research on public texts as of September 28, 2026. Each rule binds only certain licensees, and most texts leave the start of the clock to judgment. Arabic prevails for on-shore texts in every GCC state. The DIFC, ADGM and QFC publish their rules in English. A group needs its own counsel in each jurisdiction to map its entities to its duties.

## Three Gulf central banks want a first report within the hour


Across ten financial regulators, the shortest fixed first notice is one hour at the Central Bank of Bahrain (CBB), at the Qatar Central Bank for critical incidents and at the Central Bank of Kuwait for high-severity ones, under a framework launched on December 3, 2025.<sup>91,92,93,223</sup> It is four hours at the CBUAE, six hours at the Central Bank of Oman (CBO)<sup>94,217</sup>, 24 hours at the ADGM's Financial Services Regulatory Authority<sup>224</sup>, and 72 hours at the Dubai Financial Services Authority and the New York Department of Financial Services (NYDFS).<sup>225,204</sup> The US federal banking agencies allow 36 hours from the bank's determination.<sup>135,136</sup>

## 1 hour

for a first report of a serious cyber incident to the central bank in Bahrain, Qatar and Kuwait. US bank regulators allow 36 hours from the bank's determination.

Bendings analysis of regulator texts <sup>91,135</sup>

Figure 21 Three Gulf central banks expect a first report within one hour; US bank regulators allow 36 hours



Clocks start on different events: occurrence or detection, which the text does not rank (Bahrain), identification or discovery (Kuwait, Qatar, Oman's six-hour notice), not stated (UAE), awareness (ADGM, DIFC), a determination (US, NYDFS) or classification (DORA, capped at 24 hours from awareness). SAMA requires notice 'immediately' with no hour figure and is omitted. Kuwait's one hour applies to high-severity incidents. The NYDFS clock is 24 hours for an extortion payment. UK FCA 24-hour expectation is guidance and starts 2027-03-18.


Source: Regulator rulebooks and regulations retrieved 2026-09-28. Bendings analysis. n = 10

The hours are not measured from the same moment. Bahrain's call runs from occurrence or detection, Kuwait's from discovery and Qatar's from identification, while the US clock waits for a determination the bank makes and records.<sup>91,93,92,135</sup> How far that moves hour 0 is the subject of the section on start events below.

## In a scenario we built, the Gulf deadlines come first

We built a hypothetical group and ran one incident through its duties. The group is a US-listed bank holding company with a national bank, an entity licensed by NYDFS, and banks licensed in Saudi Arabia, onshore in the UAE and in Bahrain. All of them share one platform. Encryption is detected there at hour 0, a Monday. At hour 1 each Gulf bank classifies the incident as high or critical. At hour 4 US management determines that it is a notification incident for the banking agencies and a cybersecurity incident for NYDFS. At hour 30 forensics confirm that customer personal data was taken from the US, Saudi, UAE and Bahrain books. Every assumption is written down in the scenario file, and the script that computes the clocks reruns from the saved texts.

Figure 22 In one ransomware scenario the first deadline falls in Bahrain at hour 1 and the first US deadline at hour 40



Each bar runs from the event that starts a clock to its deadline. Bars cover the clocks with a fixed deadline, including conditional and pending ones. Clocks that say "immediately" and carry no hour are not drawn.

Hypothetical group: US-listed bank holding company with a national bank, a NYDFS-licensed entity and banks licensed in Saudi Arabia, the UAE (onshore) and Bahrain. Hour 0 is detection of encryption on a Monday. Trigger times are assumptions stated in scenario.csv. CIRCIA rows are not in force. Payment rows apply only if a ransom is paid. Not legal advice.

Source: Bendings analysis of regulator texts retrieved 2026-09-28 (see report/research/clocks.md, Sources). Hypothetical scenario. n = 33

Sixteen of the base-case clocks carry a fixed hour. Four fall due within 24 hours of detection, and all four are in the Gulf: the CBB phone call at hour 1 and its Section A report at hour 2, then the CBUAE first notice at hour 4 and its summary at hour 24.<sup>91,94</sup> Seven fall due within 72 hours. The

first US deadlines are the Federal Reserve and OCC notices at hour 40, 36 hours after the determination we assumed at hour 4.<sup>135,136</sup> Another 10 base-case clocks say "immediately", "as soon as possible", "without undue de-

lay" or "after resuming operations", and we did not convert those words into hours.

By hour 40 the Bahrain bank has phoned the CBB, filed a written report and called Bahrain's national cyber hotline, as a July 2025 CBB letter now directs.<sup>226</sup> The UAE bank has filed twice. The Saudi bank has notified the Saudi Central Bank (SAMA) "immediately" under two separate frameworks.<sup>227,228</sup> The data protection clocks run from hour 30, and the Saudi Data and AI Authority (SDAIA) and Bahrain's data protection authority are due at hour 102.<sup>229,230</sup> Our reading is that a group that starts drafting its first notices in New York after detection will miss the Bahrain and UAE deadlines.


The scenario lists 33 clocks. Two are CIRCIA duties not yet in force, and five apply only under a stated condition, such as paying the ransom. Of the 26 that remain, 22 start within the first 30 hours, at 13 authorities in four countries. That count is a product of our choices. It includes clocks with no fixed hour, counts SAMA's two frameworks separately and leaves out entities the group does not have. A group with a London or DIFC entity would add duties to the UK data protection regulator, the DFSA and the DIFC Commissioner, and from March 18, 2027 to the Financial Conduct Authority, all coded in our data. The order of the deadlines depends less on these choices than the count does, because the Gulf central bank clocks are the shortest in the texts.

---

## Seven kinds of event start the clocks

We coded the start event of all 96 duties. Awareness or belief starts 34, detection, identification or discovery 18, a determination or classification 12 and occurrence 10. The rest start at recovery, an earlier report or a ransom payment, or do not say. Of the nine that do not say, five are in the scenario: the CBUAE's four-hour, 24-hour and 72-hour notices, its notice of a significant personal data breach, and SAMA's five-day detailed report.<sup>94,231,228</sup> For those we chose a start ourselves.

Figure 23 Notification clocks start on 7 different kinds of event, and 9 of 96 duties do not say when the clock starts



Each row is one duty (initial notices, follow-up reports and individual notices counted separately). Trigger coding is ours; see trigger\_type and trigger in clocks.csv.

Source: Bendings coding of 96 notification duties in regulator texts retrieved 2026-09-28. n = 96

The families split by region. Seven of the 20 US duties start on a determination, which the organization makes and records. Ten of the 66 GCC duties start at occurrence, alone or as an alternative to detection, among them the CBB's one-hour, two-hour and ten-day reports and the CBO's reports by severity.<sup>91,232,217</sup> None of the 20 US duties does.

# 9 of 96

notification duties in US, GCC, EU and UK texts do not say what event starts the clock.

Bendings coding of regulator texts retrieved September 28, 2026

Where a text offers occurrence or detection, as the CBB's one-hour call does, we read the later of the two as governing. Some duties run from occurrence alone, among them the CBB's ten-day Section B report and the CBO's 12, 24 and 48-hour reports by severity.<sup>232,217</sup> Our reading is that an intrusion found a week after it began starts those clocks with most of the time gone. Seven of the ten days on the Section B report would already have run. Regulators could settle this by saying in guidance that such clocks run from detection when the occurrence could not have been detected.

## An affiliate's incident can start the New York clock

The NYDFS rule reaches across the seam on its own terms. Its definition of a cybersecurity incident includes an event that "impacts the covered entity and requires the covered entity to notify any government body".<sup>204</sup>

In no event later than 72 hours after determining that a cybersecurity incident has occurred at the covered entity, its affiliates, or a third-party service provider.

23 NYCRR 500.17(a)(1)<sup>204</sup>

So the Gulf row and the New York row are linked. Whether a Saudi or Bahraini affiliate's notice to its own central bank "impacts the covered entity" in New York is a judgment the text leaves open. We read it as a reason to treat any Gulf notice as a prompt for a New York decision, recorded with its time. Paying the ransom adds two more New York clocks, a notice within 24 hours and a written justification within 30 days.<sup>204</sup> FinCEN expects banks to consider a suspicious activity report for ransom payments, on a 30-day clock from initial detection.<sup>233,234</sup>

## The rules moved in 2025 and 2026

Several of these duties are new since 2025, and one more is pending.

- **SEC Regulation S-P.** Brokers, dealers and advisers must notify affected individuals within 30 days of becoming aware of unauthorized access to customer information, in force for all firms since June 3, 2026.<sup>235,236</sup> In a notice effective June 17, 2025, the SEC withdrew its separate proposals for broker-dealer and adviser cyber incident reporting.<sup>237</sup>
- **ADGM.** The FSRA's 24-hour cyber incident notice has applied since January 31, 2026.<sup>224,238</sup>
- **Bahrain and California.** The CBB directed banks in July 2025 to report suspicious cyber activity directly to the national cyber center.<sup>226</sup> California now sets 30 calendar days for consumer breach notices.<sup>239</sup>

One clock has not started. The Cyber Incident Reporting for Critical Infrastructure Act of 2022 sets 72 hours from reasonable belief and 24 hours from a ransom payment, but only once a final rule takes effect.<sup>240</sup> The 2026 Unified Agenda lists that rule for September 2026, a year after the statutory deadline, and the Federal Register showed none through September 25, 2026.<sup>241,242,243</sup> The scenario draws both CIRCIA clocks hollow.

## Where the data disagrees

Data protection is where jurisdictions look most alike. Saudi Arabia, Bahrain, Oman, the ADGM and the Qatar Financial Centre all set 72 hours, the same as the GDPR.<sup>229,230,244,245,246,247</sup> Qatar's onshore 72 hours comes from official guidance, not the statute.<sup>248</sup> The UAE federal law is the gap. It leaves the period to executive regulations we could not find on official portals, and it excludes banking and credit data that other legislation protects.<sup>249</sup> Here the seam is narrower than the thesis would predict.

**Figure 24 Gulf data protection laws converge on 72 hours, but the UAE federal law has no period in force**

| Authority                      | Clock (hours) | Clock as written                                                                                                                     | Starts on                 |
|--------------------------------|---------------|--------------------------------------------------------------------------------------------------------------------------------------|---------------------------|
| Kuwait (CITRA, licensees only) | 24            | within a period not exceeding 24 hours after knowledge                                                                               | knowledge                 |
| Saudi Arabia (SDAIA)           | 72            | within a period not exceeding 72 hours of becoming aware; missing information 'as soon as possible' with justification for the delay | awareness                 |
| Bahrain (PDP Authority)        | 72            | within a period not exceeding (72) hours from the date of its discovery                                                              | discovery                 |
| Qatar (PDPPL guideline)        | 72            | without delay and within 72 hours of becoming aware (guideline); the law itself sets no time limit                                   | awareness                 |
| Oman (MTCIT)                   | 72            | within (72) seventy-two hours after becoming aware                                                                                   | awareness                 |
| ADGM (Commissioner)            | 72            | without undue delay and, where feasible, not later than 72 hours; reasons for delay if later; information may be phased              | awareness                 |
| QFC (Data Protection Office)   | 72            | without undue delay and, where feasible, not later than 72 hours after having become aware                                           | awareness                 |
| EU GDPR (benchmark)            | 72            | without undue delay and, where feasible, not later than 72 hours after becoming aware                                                | awareness                 |
| UK GDPR (benchmark)            | 72            | without undue delay and, where feasible, not later than 72 hours after becoming aware                                                | awareness                 |
| DIFC (Commissioner)            |               | as soon as practicable in the circumstances; information may be provided in phases                                                   | breach (start not stated) |
| UAE federal PDPL (Data Office) |               | 'at the time it becomes aware', within the period set by the Executive Regulations (not issued/located)                              | awareness                 |
| Texas Attorney General         | 720           | as soon as practicably possible and no later than 30 days after discovery, if 250 or more Texans are affected                        | discovery                 |

Qatar's 72 hours comes from official guidance, not the statute. Kuwait's CITRA regulation binds CITRA-licensed service providers, not banks. The UAE federal PDPL leaves the period to executive regulations that were not located, and excludes banking and credit data governed by other legislation (Art. 2(2)(f)). DIFC requires notice 'as soon as practicable'. US states set days, not hours; Texas is one example.

Source: Data protection laws, regulations and official guidance retrieved 2026-09-28. Bendings analysis. n = 12

The scenario also flatters our argument in ways a reader should discount. It ignores time zones, weekends that differ across the six states, and public holidays, except in the SEC's business-day count. It assumes regulators neither extend nor waive a deadline, and supervisors can do both. Several Gulf clocks sit in frameworks that say "should" rather than "must", such as the SAMA Cyber Security Framework, and we recorded the wording rather than guess how strictly it is enforced.<sup>227</sup> Commentary cites SAMA deadlines of 2, 6, 24 or 72 hours. We could not trace any of them to an official text, so they are left out.

Finally, a short clock is not the same as a hard one. The CBB's one-hour call asks for a phone call, not a finished analysis. The 36-hour US notice asks for notice "as soon as possible", which may mean earlier than hour 40 in practice. Our figures compare the latest compliant time, which is what the texts fix, not what supervisors expect to see.

## What to do

- **General counsel:** build a clock map for each legal entity: every duty, authority, start event as written, deadline and source text, reviewed each quarter.

- **CISO:** name one notification owner for the group, with a deputy in each Gulf jurisdiction, who holds the clock map during an incident and signs off each notice.
- **Chief risk officer:** agree in writing, before any incident, which internal events count as detection, awareness and determination for each entity, and who may declare each.
- **Head of incident response:** keep pre-drafted first notices for every clock due within 24 hours, so the hour 1 and hour 2 notices in Bahrain are filled in, not written.
- **Board risk committee:** require a yearly exercise that runs the first 72 hours of one incident across every jurisdiction, and report which deadlines were missed.
- **CBB and CBO:** say in guidance that a clock written to run from occurrence runs from detection when the occurrence could not have been detected.



# 11

## Banks guard mail but rarely say where to report

We measured two seams that anyone can see from outside at 159 US and Gulf regulators and banks. The email seam is mostly closed. The reporting seam, where an outsider with a flaw to report looks for the right person, is mostly unmarked.

|                                                           |    |
|-----------------------------------------------------------|----|
| The email seam is mostly closed at the exact domain       | 80 |
| Few banks say where to report a flaw                      | 81 |
| The wall that stops scanners also stops finders           | 82 |
| Saudi Arabia and Oman stand out on DNS and mail transport | 83 |
| How we measured, and what it cannot show                  | 84 |
| Where the data disagrees                                  | 84 |

Between 17:02 and 17:14 UTC on September 28, 2026, we sent DNS queries and at most three web requests to each of 159 financial institutions: 12 US authorities, 50 large US banks, 18 Gulf authorities and 79 Gulf banks. In those twelve minutes, 147 of them told the world's mail servers to quarantine or reject mail that falsely claims to come from their domain. Seven served a security.txt file telling a stranger where to report a flaw. Three of those files were valid.

Everything in this chapter is our analysis of public DNS records and public web responses, which anyone can check in minutes without access or permission. We report groups only and never name an institution.


### Two outside seams

The email seam is where a receiving mail server decides whether a message that claims your domain really came from you. The reporting seam is where someone outside, who has found a flaw in your systems, decides whether and how to tell you.

## The email seam is mostly closed at the exact domain

DMARC lets a domain owner tell receivers what to do with mail that fails authentication: nothing (p=none), put it aside (p=quarantine) or refuse it (p=reject).<sup>250</sup> We count an institution as enforcing when its policy is quarantine or reject and applies to all mail. By that test, 147 of 159 enforce, or 92 percent.

Figure 25 147 of 159 measured regulators and banks tell mail receivers to quarantine or reject mail that fails DMARC



Each bar is one group. The dark segments are the institutions a receiver is told to protect in full.

DMARC record at \_dmarc.<primary domain>, with fallback to the organizational domain. Counts: US authorities: reject 12, quarantine 0, partial 0, none 0, missing 0; GCC authorities: reject 10, quarantine 5, partial 1, none 2, missing 0; US banks: reject 44, quarantine 3, partial 0, none 2, missing 1; GCC banks: reject 61, quarantine 12, partial 3, none 2, missing 1. Measures the primary domain only; lookalike domains, subsidiaries and other sending domains are not covered.

Source: Bendings measurement of DNS (Cloudflare and Google validating DNS-over-HTTPS resolvers) and at most three HTTPS GET requests per institution, 2026-09-28T17:02:38Z to 2026-09-28T17:14:13Z UTC. Targets: FFIEC NIC large holding companies (2026-06-30), FSOC member agencies, NYDFS, FINRA, CISA, and GCC central bank registers. Bendings analysis. n = 159

All 12 US authorities publish p=reject. Forty-seven of the 50 US banks enforce, 44 of them at reject. Seventy-three of the 79 Gulf banks enforce, 61 at reject. Gulf authorities trail at 15 of 18. Of the 12 that do not enforce, four apply the policy to only part of their mail, six are at p=none and two have no valid record. All 159 publish an SPF record, the list of servers allowed to send for the domain.<sup>251</sup>

## 67 of 79

Gulf banks end their SPF record in the strict "-all", telling receivers to fail mail from unlisted servers. Among the 50 large US banks, 27 do.

Bendings measurement, September 28, 2026

On the SPF record's final instruction, Gulf banks are stricter than large US banks, 67 of 79 against 27 of 50 ending in "-all". The rest of both groups end in the softer "~all", apart from one Gulf bank with no "all" term. Strictness has a cost. Three of the 79 Gulf bank records need more than the 10 DNS lookups RFC 7208 allows, which breaks their evaluation, and four more sit exactly at the limit. No US bank record went over 10, but 23 of the 50 use macros we cannot expand, so those counts are lower bounds.

An enforced DMARC policy covers mail that uses the institution's exact domain in the From line. It does nothing about a lookalike domain, a subsidiary or brand domain with a weaker record, a supplier that sends invoices for the bank, or a display name that reads like the bank over an unrelated address. A phishing path that starts on a lookalike domain and runs through an employee to a credential and on into the network, **H / I / N**, never touches the domain we measured. Chapter 4 follows those paths.

---

## Few banks say where to report a flaw

A customer or a researcher who finds a flaw at a bank needs an address to report it. RFC 9116 gives one fixed place to look, a small text file at a standard address on the website that names a contact and an expiry date.<sup>252</sup> It costs nothing to publish. We asked each institution's main web host for it.



**When security vulnerabilities are discovered by researchers, proper reporting channels are often lacking. As a result, vulnerabilities may be left unreported.**

RFC 9116, the security.txt standard<sup>252</sup>

**Figure 26 None of 79 GCC banks and 5 of 50 large US banks published a security.txt file on their main website**



A file counts only if it is a text file with at least one RFC 9116 field. It counts as valid only with a contact, a single expiry date in the required format that has not passed, and delivery over HTTPS.

One GET to `https://<web host>/.well-known/security.txt`, one redirect followed when the request budget allowed. HTML pages, soft 404s and block pages do not count as a file. Found / valid: US authorities 1/12 / 1/12; GCC authorities 1/18 / 0/18; US banks 5/50 / 2/50; GCC banks 0/79 / 0/79. Bot protection, TLS errors or timeouts prevented a clean read for US authorities 1; GCC authorities 5; US banks 11; GCC banks 18. Institutions may publish a disclosure policy elsewhere; this tests only the RFC 9116 location on the main web host.

Source: Bendings measurement of DNS (Cloudflare and Google validating DNS-over-HTTPS resolvers) and at most three HTTPS GET requests per institution, 2026-09-28T17:02:38Z to 2026-09-28T17:14:13Z UTC. Targets: FFIEC NIC large holding companies (2026-06-30), FSOC member agencies, NYDFS, FINRA, CISA, and GCC central bank registers. Bendings analysis. n = 159

None of the 79 Gulf banks served a file. For 18 of them bot protection, a TLS error or a timeout stopped a clean read, so the rate among the 61 we could read is still zero. Five of 50 large US banks served one, and two of the five were valid. One had expired, one had no expiry date and one gave a date without a time. One of 12 US authorities served a valid file. One of 18 Gulf authorities served a file, and it had no expiry date.

At 16 of the 79 Gulf banks, the request for the file returned an ordinary web page instead. A person sees something that looks like an answer; a program gets nothing it can use.

## 0 of 61

Gulf banks whose websites we could read served a security.txt file. Five of 50 large US banks did, and two of those files were valid.

Bendings measurement, September 28, 2026

The US authority numbers need context. CISA's Binding Operational Directive 20-01, issued September 2, 2020, requires federal civilian agencies to publish a vulnerability disclosure policy at a fixed path on their websites and to list a security contact in the .gov registry. It says of security.txt that "use is not required under the directive".<sup>253</sup> All 11 US authority .gov domains in our list carry a security contact in

CISA's registry data.<sup>254</sup> We did not request the disclosure policy pages, so we cannot say how many publish one.

The directive also names the finder's fear. Its guidance notes that the federal government has a reputation in the security community for being "defensive or litigious", and the directive requires agencies to commit not to pursue legal action against good faith research.<sup>253</sup> We found no Gulf registry field of this kind to check.

## The wall that stops scanners also stops finders

Our requests named their sender and gave a contact address. Bot protection turned them away anyway. The home page request met a block page or HTTP 403 at 8 of 50 US banks and 8 of 79 Gulf banks. The security.txt request was blocked at 7 of 50 US banks and 15 of 79 Gulf banks, 22 of 129 banks in all.

We did not retry as a browser. All requests came from one network location in the United States. If any of these sites filter by the requester's country, part of the gap between the regions reflects where we stood. Our judgment is that many tools researchers use to check for a disclosure file look to a bot filter much like ours. A control built to keep automated traffic out also keeps out the automated check for how to re-

port a flaw. The two decisions usually belong to different teams, which is what makes it a seam.

# Saudi Arabia and Oman stand out on DNS and mail transport

The figure sets Gulf banks by country beside large US banks on every record we checked. Between 86 and 100 percent of banks in each Gulf state enforce DMARC, against 94 percent of US banks.

Figure 27 Gulf banks match large US banks on DMARC enforcement and use strict SPF more often, but none publish security.txt

|                     | DMARC quarantine or reject at | SPF ends in -all | DNSSEC signed and validated | MTA-STS in enforce mode | TLS-RPT record | CAA record | security.txt found | HSTS header |
|---------------------|-------------------------------|------------------|-----------------------------|-------------------------|----------------|------------|--------------------|-------------|
| Saudi Arabia (n=14) | 86%                           | 79%              | 64%                         |                         |                | 7%         |                    | 79%         |
| UAE (n=24)          | 92%                           | 88%              | 17%                         |                         | 4%             | 8%         |                    | 92%         |
| Qatar (n=9)         | 89%                           | 100%             |                             |                         |                | 11%        |                    | 56%         |
| Kuwait (n=10)       | 100%                          | 90%              | 10%                         |                         |                |            |                    | 80%         |
| Bahrain (n=12)      | 100%                          | 83%              | 25%                         |                         |                |            |                    | 83%         |
| Oman (n=10)         | 90%                           | 70%              |                             | 40%                     | 70%            | 10%        |                    | 70%         |
| US banks (n=50)     | 94%                           | 54%              | 22%                         |                         | 2%             | 28%        | 10%                | 76%         |

Denominators are small. Read the Qatar, Kuwait, Oman and Bahrain rows as counts of 9 to 12 banks, not as rates.

Each cell is the share of banks in the column group. Small denominators: Qatar 9, Kuwait 10, Oman 10, Bahrain 12. HSTS here uses all institutions as the denominator, including hosts that blocked or failed; figure posture-hsts uses reachable hosts only. GCC bank lists: locally incorporated banks on each central bank register (Bahrain: retail licences only).

Source: Bendings measurement of DNS (Cloudflare and Google validating DNS-over-HTTPS resolvers) and at most three HTTPS GET requests per institution, 2026-09-28T17:02:38Z to 2026-09-28T17:14:13Z UTC. Targets: FFIEC NIC large holding companies (2026-06-30), FSOC member agencies, NYDFS, FINRA, CISA, and GCC central bank registers. Bendings analysis. n = 129

DNSSEC signs a domain's DNS answers so that a resolver can detect a forged one. We count a domain as validated when its parent zone holds a DS record and two public validating resolvers mark the answer authenticated. Nine of 14 Saudi banks validate. Elsewhere in the Gulf, 8 of 65 banks do, three of 12 in Bahrain, four of 24 in the UAE, one of 10 in Kuwait and none in Qatar or Oman. Only 2 of 18 Gulf authorities validate. In the United States, 9 of 12 authorities validate, against 11 of 50 banks. Eleven more domains sign their zones but never published a DS record at the parent, which leaves the signatures unused.

MTA-STS lets a domain tell sending servers to deliver mail only over authenticated TLS.<sup>255</sup> Five of 159 institutions publish a policy in enforce mode, all in Oman, four banks and one authority. None of the 69 Gulf banks outside Oman publishes one, and none of the 62 US authorities and banks does. The figure also shows HSTS, which most banks in both regions send.

## How we measured, and what it cannot show

The lists come from public registers. US banks are the top 50 holding companies by assets on the FFIEC National Information Center list as of June 30, 2026, with one substitution for a completed merger.<sup>256</sup> US authorities are the nine federal agencies at the July 15, 2026 meeting of the Financial Stability Oversight Council plus NYDFS, FINRA and CISA.<sup>257</sup> Gulf banks are the locally incorporated banks on each central bank's register; Gulf authorities are each state's central bank, markets regulators and national cybersecurity authority.<sup>258,259,260,261,262,263</sup> Four institutions whose domains we could not confirm are left out. The US list holds only the largest banks and the Gulf list holds every local bank, large and small, so bank comparisons between the regions are not like for like.

Each institution is one web domain and, for nine of them, one separate published mail domain. The results are a single point in time. DMARC is an instruction; its effect depends on receivers that apply it. We did not test DKIM, which cannot be discovered passively.

## Where the data disagrees

This report argues that intrusions cross seams that organizations leave unexamined. The email seam at the exact domain is not one of them. Ninety-two percent of these institutions enforce DMARC, and Gulf banks match large US banks on enforcement and beat them on strict SPF.

The reporting seam result is weaker than it looks. security.txt is one channel among several. A bank may run a bug bounty, publish a disclosure page elsewhere or answer reports sent to a general security address, and we tested only the standard location on one host. Absence means "not at the standard place". It does not mean "unreachable". Our US agency figures would have told the wrong story without the .gov registry. We did not check the reporting pages of the Gulf national CERTs, and the Gulf may have channels we have no public way to see.

Finally, the measurement cannot link posture to harm. We have no evidence that institutions without a security.txt file received fewer reports or suffered more incidents. The

claim is narrower. The path from finder to defender is unmarked at the one place built to mark it.

## What to do

- **CISOs of Gulf and US banks:** publish a valid security.txt file on the main web host this quarter, with a monitored contact, an expiry date in RFC 3339 format and a link to a disclosure policy, and put its renewal on the same calendar as your certificates.
- **Heads of web and bot defense:** exempt /.well-known/security.txt and your disclosure policy page from bot challenges, and test from a plain HTTP client that both return the file.
- **Heads of email security:** list every domain that sends mail as the institution, including subsidiaries, brands and suppliers, bring each to p=reject, and check that no SPF record exceeds 10 lookups.
- **Heads of infrastructure:** sign the primary domain with DNSSEC and publish the DS record at the parent, and add MTA-STS and TLS reporting for inbound mail.
- **Gulf regulators and national cybersecurity authorities:** set an expectation, as BOD 20-01 did for US federal agencies, that supervised banks publish a disclosure policy and a contact. Agree with the public prosecutor or national cyber authority in each state what protection a good faith reporter has, because a bank's own promise does not bind the state.



# 12

## Test the crossings, not only the systems

Keep a register of seams, measure how many have ever been crossed in a test, design scenarios around bends, and arrange the authority to cross them before anyone needs it.

|                                                            |    |
|------------------------------------------------------------|----|
| Keep a seam register                                       | 86 |
| Five measures a board can ask for                          | 86 |
| Design each scenario around a bend                         | 87 |
| Arrange the authority before it is needed                  | 87 |
| Regulators can require the crossings their rules leave out | 88 |
| Where this method is weak                                  | 88 |

Most testing programs are planned from an asset inventory. List the applications, networks, buildings and people, then schedule a test for each. An inventory lists systems. It has no rows for the crossings between them, so tests planned from it never cross one.

Chapter 2 found that tests cover the steps of documented intrusions but miss the joins, and that most paths which left a red team's scope ran through another organization's systems or people. We propose seam-led testing, which plans from the bends an adversary would need and asks which of them the organization has ever seen crossed end to end.

# 70 of 78

documented intrusion paths had every known step on the victim's own systems in a domain some single-domain test covers. A typical red team scope could follow only 33 of the 78 from first step to last.

Bendings incident corpus, 2013 to September 2026, chapter 2

| Field                      | What goes in it                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Crossing                   | From and to, in bend notation. H / I is the service desk resetting a credential; I <sub>s</sub>   N is a supplier's account reaching your network. |
| Public example             | A documented case in which an adversary made this crossing, so the row is argued from evidence and not imagination.                                |
| Owner on each side         | A named person, not a team, for each side of the seam.                                                                                             |
| Control at the crossing    | What should stop the crossing, and who operates it.                                                                                                |
| Who would see it           | The log that records the crossing, where that log lives, and the team that reads it.                                                               |
| What stops on the far side | The operations that halt if the near side is lost, including those you would halt yourselves as a precaution.                                      |
| Last crossed in a test     | Date and reference of the last test that crossed it end to end, or "never".                                                                        |
| Authority needed           | Whose permission a test of this crossing requires: a supplier, a landlord, a second regulator, a second country.                                   |
| Jurisdiction on each side  | The country and regulator on each side, where they differ. Bend notation does not mark this change.                                                |

## Five measures a board can ask for

Boards usually see the number of findings, the share closed and the time to close them. Those numbers say nothing about the crossings between the systems tested. Five measures do.

- 1 **Seam coverage.** The share of registered crossings exercised end to end in the last twelve months, shown

## Keep a seam register

A seam register is a short list, one row per crossing, kept beside the risk register and owned by the same person. It need not be complete on the first day. Twenty rows that are true are worth more than two hundred that are guessed. Appendix C has a blank register and starter rows drawn from bends that recur in the corpus.

beside the number of rows in the register.

- 2 **Ownerless seams.** The number of registered crossings without a named owner on both sides. The target is zero, and zero is reachable.
- 3 **Crossing detection.** In tests, the share of crossings that monitoring saw within a time agreed in advance, reported per crossing and not as an average.
- 4 **Supplier crossings tested.** The number of crossings into or out of a supplier that were tested with the supplier taking part.

- 5 **Longest path tested.** The most bends exercised in a single scenario in the last twelve months, reported with the crossings it included. A program whose longest tested path has one bend has not tested how its controls hold together.

All five come from test reports and the register, with no new tooling. Each can be gamed. A short register makes coverage look high, and one that leaves out hard crossings has no ownerless seams. So the register needs a floor. It should include every starter row in Appendix C that applies, plus a row for every supplier whose people or credentials reach the organization's systems. Show the board the row count beside the coverage share and the crossings beside the longest path, so it can judge whether they were the ones that matter. We publish no benchmark for these measures, because no public data could set one.

## Design each scenario around a bend

A seam-led scenario names the crossing first and the systems second. Eight scenarios cover bends found in the public record, from the most common to a few that are rare there. Each has a stopping point agreed before it starts.

- **Service desk to administrator, H / I / C.** A caller asks the service desk, including any outsourced desk, to reset a privileged user's MFA. The test ends at the first administrative action.
- **Supplier access to the core, I<sub>s</sub> | N.** How far a supplier's remote-support credentials reach, with the supplier taking part in writing.
- **Appliance to identity, E / I / C.** Assume the VPN or firewall is compromised, and see what its service accounts, stored credentials and management plane reach.
- **Integration token to data, C<sub>s</sub> | C.** What a token held by a connected SaaS application can read in your tenant.
- **Lobby to directory, P / N / I.** A device left in a meeting room or plugged into a public port, and the path from there to a directory account.
- **Approval to payment, I / V.** From an approver's account to a payment release, stopping at the last point before money moves.

- **Engineering workstation to controller, N / O.** In a lab or on a replica, never on live safety systems.
- **Assistant to action, A / C.** Content planted where an AI assistant with tool access will read it, and what the assistant can then be made to do.

### One scenario, written out

A bank in the Gulf outsources its first-line service desk to a provider in another country and runs its payment approval workflow in a cloud tenant. The scenario is H<sub>s</sub> | I / C / V. A caller persuades the provider's agent to reset an approver's MFA, the tester signs in as the approver, reaches the workflow and stops before a payment is released. To run it, the bank needs the provider's written participation, a control group at the provider who can stop the call, confirmation that recording the call is lawful where the agent sits, a named person at the bank who can halt the test, and a stopping point in the payment system agreed with its owner. The test result is four facts: whether the reset happened, whether anyone at either firm saw it, how long the approver's session lasted before someone noticed, and whether the payment workflow's own controls would have held.

## Arrange the authority before it is needed

A test can cross only the seams that every owner along the path has authorized. That rule holds. What can change is how early the authority is arranged.

- **Write the right to test into supplier contracts,** including participation in joint scenarios and access to the supplier's logs for the test window.
- **Get letters from each owner of a physical site,** not only the tenant.
- **Name the law of each jurisdiction a scenario touches** in the rules of engagement, and who checked it. Some Gulf rules also limit where testers may sit (chapter 9).
- **Record the paths that authority stopped.** When a test reaches a seam it is not authorized to cross, the tester should stop, record the path and report it. Those recorded paths are the first rows of the next register.

## Regulators can require the crossings their rules leave out

Chapter 9 counts the gaps. Of 41 US and Gulf testing regimes, 1 requires a physical intrusion test, 5 a test of people and 9 put supplier-run systems in scope, and none of the 50 regimes we read sets rules for AI-driven testing. Our four proposals are each aimed at a body that can adopt them. Only the second reaches past the regulated firm, and it does so through the firm's own contracts.

- 1 **SAMA and the Central Bank of Kuwait**, the two Gulf central banks with a fixed red team cycle: in the next revisions of FEER and CORF, require one scenario per cycle that crosses at least two seams in one path, recorded in bend notation. Until then, a bank can propose one, subject in Saudi Arabia to the approval of SAMA's green team.<sup>132,93</sup>
- 2 **US and Gulf regulators, through outsourcing and cloud rules**: require contracts with shared providers to include taking part in a pooled test led by their regulated customers, with at least one scenario on the provider's own systems, as DORA does.<sup>89,90</sup> A central bank can then coordinate the pool without supervising the provider.
- 3 **Home and host supervisors of the same banking group**: agree bilaterally that one threat-led test of a shared platform counts in both states, once tester location and accreditation rules are reconciled. Today SAMA's green team approves the provider, Saudi Arabia's critical systems controls bar remote access from outside the Kingdom, and the Qatar Central Bank requires an external provider "with regional offices", so no single test team can satisfy every state.<sup>132,130,201</sup>
- 4 **Regulators that set testing rules**: set minimum controls for AI systems that test, along the lines of the profile in chapter 8.

place of a test. Seam-led testing is itself untested, built from the crossings in the corpus with no public evidence that it finds more than a conventional red team of the same cost. Each organization will learn from its own test reports whether it does, and the five measures would show the difference. It also costs more per scenario than a single-domain test, because it needs more permissions and more people. It does not replace penetration testing. It decides where the next one should start.

## What to do

- **CISO**: start a seam register this quarter at the floor set out above, and name an owner on both sides of each row.
- **Head of offensive security**: redesign next year's red team around two registered crossings that have never been tested, including one owner seam.
- **Procurement and supplier risk**: add test participation and log access to the renewal terms of every supplier whose people or credentials reach your systems.
- **Board risk committee**: ask for the five measures twice a year, with the register's row count beside them, starting with ownerless seams.
- **SAMA and the CBK**: invite one seam-crossing scenario in the next FEER or CORF test while the frameworks are revised.

## Where this method is weak

A register of crossings is only as good as the people who write it, which is why every row should cite a public case. Some seams cannot be tested without a partner who will not agree, and some should never be tested on live systems, such as safety controllers. For those, the honest entry is "never crossed, and why", with a tabletop exercise in

# How Bendings works at the seams

The one place in this report where we describe ourselves and what we sell. Everything here is a standing commitment, not a claim about results.

Bendings has worked for governments, banks and asset managers in the United States and the GCC states since 2013, by referral. We run reconnaissance, red team, penetration testing, physical and social-engineering testing, blue team support and incident response, and we test operational technology, embedded devices and radio interfaces in our own hardware lab as well as on site. That is the testing this report finds missing, so we have an interest in its proposals.

## A rule about paths we may not follow

Systems, sites and people outside an engagement's scope are never tested, even when a path to them opens during the work. We record the path and report it, as our published method states. Each recorded path is a crossing the client now knows exists and can decide to authorize, fix or accept.

## Bendings Frame

Frame is the operating system our team works in. It brings logs, alerts, network and device data from many sources into one live picture as they arrive, with cameras, access control and sensors in the same picture as the networks. It runs inside a client's network, in-country, fully air-gapped or hosted by us. Only Bendings operators work in it. Clients see what it finds.

## Bendings Gauge

Gauge is our own AI model for continuous red and blue team work. It runs on Frame, inside limits each client ap-

proves. We publish five commitments for it:

- 1 **Scope it cannot leave.** Gauge itself refuses any target outside the signed scope, whatever it finds along the way.
- 2 **Operators approve risky steps.** Exploitation, lateral movement and other risky actions wait for a Bendings operator to approve them.
- 3 **You can stop it at any time.** A named person on the client's side can halt Gauge immediately.
- 4 **It runs where your data lives.** Inside the client's network, in the client's country or fully air-gapped. Client data stays there and is never used to train the model.
- 5 **Every action on record.** Gauge logs each action it takes, and the complete log is handed to the client.

Chapter 8 proposes ten controls for any AI system used in testing. Our five commitments correspond, at most partly, to controls 1, 3, 4, 7 and 8. Control 1 asks for scope enforced outside the model, where ours says Gauge itself refuses. Control 4 asks for a stop that works without going through the provider, and control 7 for a log the agent cannot touch. Our commitments say neither. We publish nothing on controls 2, 5, 6, 9 and 10. They are egress denied by default, a monitor independent of the agent, stopping when a task cannot be done in scope, disclosure of the models and outside AI services used, and human sign-off with independent evidence. Ask us about all ten, as chapter 8 tells buyers to ask every provider.

We publish no benchmark claims for Gauge until we have results others can reproduce. Engagements begin with a conversation and a signed scope, never with a test, and client names stay private. To discuss anything in this report, start at [bendings.com/engage](https://bendings.com/engage).

## Appendix A

# Method and limitations

How each dataset was built, what it can and cannot show, and where to find the data and code.

Every number in this report is either computed by Bendings from public data or quoted from a named source. None comes from client work. The snapshot date for all data is September 28, 2026, and every source in the notes records the date we read it. Chapters 5 and 8 also describe events reported in October 2026, read from the primary sources in the notes. None of them is counted in any dataset. Where a source was a company describing its own product or incident, the notes say so. We did not use documents marked for restricted distribution, even where copies circulate publicly without the issuer's consent. The case files, derived data, figure specifications, scripts and sources are published at [bendings.com/research](https://bendings.com/research), and the report is built from them by a script. Raw measurements that would name institutions or ransomware victims are held back; the scripts show how to collect them again.

## The incident corpus

**What it is.** 78 intrusions with first access or disclosure between January 1, 2013 and September 28, 2026, each coded as an ordered path of steps. Each step records a domain, an owner, a one-sentence action and the sources that support it. Cases also record sector, region, actor type, objective, who discovered the intrusion, dwell time where both dates are published, and where harm landed.

**Inclusion.** A case needed a detailed account, from a primary source where one exists (a government advisory or indictment, a court filing, a regulator's decision, the victim's own disclosure or the responder's report), and enough detail to code at least two steps. Outages without an adversary and denial-of-service-only events were excluded, as were incidents in which AI systems under evaluation reached real organizations without anyone directing them there. Government and technology victims make up 15 cases each and finance 13, and a Gulf slice holds 15 cases. We chose the candidates and we also hold the thesis. The

research notes list the candidates considered and not coded, with the reason for each.

**Coding conventions.** Appendix B gives the rules. Two matter for the headline measures. Where the sources show that an intrusion began at a supplier or customer without saying how, the first step is recorded as undisclosed but owned by that party, so the seam stays visible. Nine paths begin this way. Some cases are campaigns against many victims, such as SolarWinds and the Exchange server intrusions of 2021, and are coded as the path a typical documented victim experienced.

**Reliability.** A second analyst coded 16 cases drawn at random (seed 20260928) from the 70 rated high or medium confidence, from the sources and without seeing the first coding. On whether the path passed through another owner, and whether it left a typical red team scope, the two agreed in 16 of 16. They agreed on the entry domain in 15 of 16 and on the exact sequence of domains in 8 of 16. Across 144 decisions on which domains each path touched, Cohen's kappa was 0.88, a figure lifted by the many easy agreements that rare domains were absent. The second analyst's notes on at least three well-known cases record steps first coded from memory and then dropped because no source supported them, so on famous cases the second coding is not fully independent. No low-confidence case was recoded.

## What the corpus measures can bear

**Bend counts follow from the coding.** A case needs at least two steps, a new step starts only when a different control had to fail, and consecutive steps with the same domain and owner are merged. Almost every step boundary is therefore a bend. The number of bends in a path measures how many separate control failures the sources document.

It is not independent evidence that intrusions bend. Better-documented cases show more, a mean of 2.48 crossings for the 31 high-confidence cases, 2.18 for the 39 medium and 1.0 for the 8 low. The report rests its argument on where bends fall and whose systems they cross.

**Test coverage.** A path counts as walkable by a test type only if every known step falls in the test's domains and no step is owned by anyone other than the victim. The domains we credit are these. External test, edge and identity; internal test, internal IT and identity; cloud review, cloud

and identity; phishing exercise, human and identity; physical test, physical; OT assessment, OT; typical threat-led red team, human, identity, edge, internal IT, cloud and value transfer. An undisclosed first step does not count against a test. The steps are well covered. In 70 of 78 paths every known step on the victim's own systems falls in a domain some single-domain test covers, and in 69 inside a typical red team scope. The joins are not. The table shows how far the whole-path counts move under other reasonable definitions.

| Measure, 78 paths unless stated                                          | Paths |
|--------------------------------------------------------------------------|-------|
| Walkable by one single-domain test, as defined above                     | 9     |
| Same, with undisclosed steps counted against the test                    | 4     |
| Same, crediting phishing and external tests with an internal foothold    | 28    |
| Of the 9, walkable only because the first step is undisclosed            | 5     |
| Walkable by a typical red team scope                                     | 33    |
| Left that scope                                                          | 45    |
| Left it through another owner's systems or people                        | 38    |
| High and medium confidence only (70): through another owner              | 36    |
| High and medium confidence only (70): left red team scope                | 43    |
| High and medium confidence only (70): walkable by one single-domain test | 7     |

A path can leave the red team scope for more than one reason. Of the 45, 38 pass through another owner, 6 reach OT, 4 need a physical step and 1 an AI system.

**A scale for noise.** The corpus is a purposive sample of well-documented cases. It favors large victims, US disclosure rules and the cases responders chose to publish, and Gulf cases rest more often on press and foreign responders. Its counts do not estimate rates for all intrusions. To show how much a count could move by chance, the data files give 95 percent Wilson intervals for the main shares. Passing through another owner, 38 of 78, runs from 38 to 60 percent. Leaving a typical red team scope, 45 of 78, runs from 47 to 68 percent. Walkable by one single-domain test, 9 of 78, runs from 6 to 21 percent. Identity in the path, 53 of 78, runs from 57 to 77 percent. With 78 cases, one case moves a share by about 1.3 points. We treat any comparison whose intervals overlap widely, such as between eras or regions, as descriptive.

## The CISA Known Exploited Vulnerabilities catalog

Catalog version 2026.09.27, 1,728 entries. The 287 entries of the launch day, November 3, 2021, are excluded from trends; "since 2022" means the 1,417 entries added from January 1, 2022 to September 27, 2026. Publication dates and severity scores come from the NVD CVE API. Each of 726 vendor and product pairs is mapped by hand-written rules to one of 19 categories and two flags, internet-facing by design and sold as a security control; 89 debatable pairs are marked and a sensitivity range is reported. The catalog lists only flaws with reliable evidence of exploitation, a CVE and a clear remediation, not all exploitation.

Year-on-year shares are headlined for recently published flaws, those published in the year of listing or the year before. In January to September 2022, 377 of 521 additions (72 percent) were flaws published two or more years earlier, as CISA loaded its backlog, which distorts every 2022

share. Days from publication to listing measure how a flaw was disclosed as much as how fast attackers moved. A flaw disclosed because it is already being exploited is listed within days. EPSS scores for the day before each listing come from FIRST's API for the 843 additions from March 7, 2023. Additions published on the day of listing or later have no such score and are left out of EPSS shares.

---

## US securities filings

All Form 8-K and 8-K/A filings with Item 1.05 from December 18, 2023 to September 28, 2026, found by EDGAR full-text search and checked against the submissions record of every filer, plus 8-Ks that disclosed a cybersecurity incident under Items 8.01 or 7.01. 185 filings, 135 incidents, 130 companies. Explicit rules coded each first disclosure, a language model coded the same variables as a second rater, every disagreement was settled by hand, and a random sample of 20 filings was hand-coded to check both. The Item 1.05 set matched 45 of 46 rows of a public law-firm tracker on filing date. Counts are of filers, not supplier events. Seven filings concern one CDK Global incident. Differences between sectors at this size, such as 10 of 28 financial filers against 26 of 107 others describing a third-party route, are descriptive and within what chance would produce.

---

## Ransomware leak-site posts

Counts come from the RansomLook tracker, with the archived ransomwatch dataset as a cross-check, because the aggregator with country fields prohibits commercial use of its data. A strict rule counts posts with a Gulf country-code domain, flag or labeled country; a broad rule adds Gulf country names, demonyms and major cities. The strict rule is a lower bound. About one in five posts the broad rule adds is a foreign firm with a Gulf office. Posts are criminal claims, not confirmed incidents, and victims who paid are absent. No victim is named in any derived file.

---

## Posture measurement

159 institutions: 12 US authorities, the 50 largest holding companies on the Federal Financial Institutions

Examination Council's list, 18 Gulf authorities and 79 locally incorporated Gulf banks from central bank registers. Measured once, from 17:02 to 17:14 UTC on September 28, 2026, from one network location in the United States, using validating public resolvers and at most three HTTPS requests per institution, one primary domain each. The US list holds the largest groups and the Gulf list every local bank, so the regions are not like for like. Results are published in aggregate only.

---

## Testing regimes and reporting clocks

50 testing regimes (12 US, 29 Gulf, 9 benchmarks) coded on 17 dimensions from the regulators' own texts, with a section reference for each coded cell and a short quotation wherever the text addresses the dimension. Two Gulf regimes have no public text. 96 notification duties (20 US, 66 Gulf, 10 benchmark) with their clock, trigger, recipient and citation; 87 of the 96 quotations were matched automatically to the saved text and the rest checked by hand. One analyst coded both sets, and no cell was coded a second time. Codes describe what texts say, not how supervisors apply them. Nothing here is legal advice.

---

## AI evidence

Chapter 8 draws on threat reports by AI developers and security firms, government evaluations, the organizations' own incident accounts, regulator texts and CVE records. Almost all of it is self-reported, and marked as such. Our own computations are a count of CVE records that name prompt injection and simple shares of totals the FBI published.

## Appendix B

# Bend notation, on one page

A reference card for writing any intrusion, test scenario or incident timeline as a path of domains and owners.

## Domains

| Code | Domain                 | Which control had to fail                                                                                            |
|------|------------------------|----------------------------------------------------------------------------------------------------------------------|
| H    | Human                  | A person's judgment. Deception, persuasion, bribery, coercion or recruitment.                                        |
| I    | Identity               | Accounts, credentials, sessions, tokens, keys, MFA, identity providers, and the processes that issue and reset them. |
| E    | Edge                   | Systems exposed to the internet on purpose, including DNS and registrar records.                                     |
| N    | Internal IT            | Internal network, endpoints, servers, directory, hypervisors, backups, build systems.                                |
| C    | Cloud and SaaS         | Cloud control planes and storage, SaaS tenants and their integrations.                                               |
| O    | Operational technology | Industrial control, safety instrumented systems, building management, controllers.                                   |
| P    | Physical               | Premises, hardware, removable media, radio proximity, devices in transit.                                            |
| V    | Value transfer         | Payment and card switches, SWIFT, ATM networks, trading and settlement, wallet signing.                              |
| A    | AI systems             | Models and agents with access to data, tools or code.                                                                |
| ?    | Not disclosed          | The sources do not say how this step happened.                                                                       |

## Owners and separators

| Mark      | Meaning                                                                                | Example              |
|-----------|----------------------------------------------------------------------------------------|----------------------|
| no letter | The step is on the victim's own systems or people                                      | I                    |
| s         | Supplier: vendor, outsourcer, managed service provider, software publisher, contractor | H <sub>s</sub>       |
| c         | Customer of the victim                                                                 | C <sub>c</sub>       |
| p         | An employee's personal device or account                                               | I <sub>p</sub>       |
| t         | Any other party: carrier, registrar, platform, neighbor                                | N <sub>t</sub>       |
| /         | Domain bend: the next step is in a different domain                                    | I / N                |
|           | Owner seam: the next step belongs to someone else                                      | C <sub>s</sub>   C   |
|           | Both at once                                                                           | H <sub>s</sub>   / I |
| .         | Same domain and owner, a separate control failed                                       | N . N                |

# Rules

- 1 Choose the domain by the control that had to fail. Phishing that harvests a password is H / I. A malicious attachment that runs code is H / N. A stolen password used on a VPN is I, not E. This rule puts every use of a credential in identity, so identity carries more bends than it would under a notation that folded credential use into the next domain.
- 2 Start a new step each time a different control had to fail. Merge repeated actions with the same domain and owner.
- 3 Stop at the objective, the step that stole, encrypted, moved money, sabotaged or disclosed.
- 4 Where the first step is unknown, write ? and keep the owner if the sources give it. If they show the intrusion

began at a supplier without saying how, write ?<sub>s</sub>, so the owner seam stays on the record.

- 5 Record consequences separately. Where harm reached a domain the adversary never acted in, note it beside the path, for example "O by decision" when an operator halted a plant as a precaution.
- 6 Record a change of jurisdiction beside the path, not in it. When the next step sits in another country or under another regulator, as with a supplier's data center abroad or a group entity supervised by a second central bank, write the jurisdiction on each side in the seam register's jurisdiction column. The notation has no mark for it, so that paths stay comparable across countries.

Bend notation records domain and owner, which MITRE ATT&CK does not. It is meant to sit beside ATT&CK. Each published case file also lists the ATT&CK technique for a step where the mapping is clear, for teams that build emulation plans from it.

# Worked examples from the corpus

| Case                      | Path                                                         |
|---------------------------|--------------------------------------------------------------|
| Bangladesh Bank, 2016     | H / N / V                                                    |
| Uber, 2022                | N <sub>p</sub>   / I / C                                     |
| 3CX, 2023                 | N <sub>s</sub>   N / I / N   N <sub>e</sub>                  |
| SEC X account, 2024       | H <sub>t</sub>   / I / C                                     |
| European Commission, 2026 | N <sub>s</sub> / I <sub>s</sub> / N <sub>s</sub>   N / C . C |

## Appendix C

# A seam register to start from

Eighteen crossings from the public record, each with a documented example, ready to be given owners and a test date.

Copy the rows that apply, delete the rest and add your own. Each row should name a person on both sides of the seam before it is considered live. The second column counts how often the crossing appears between consecutive steps in the 78-case corpus. For a domain bend, that is every change between those two domains, whoever owned the steps. For an owner seam, it is every change from that owner to the victim, in any domains. The counts

show frequency in documented cases, not risk to any one organization.

The last two columns are prefilled as a starting point. "Authority needed to test" names whose permission a test of the crossing requires. "Jurisdiction on each side" is where to record a change of country or regulator, which bend notation does not mark (Appendix B, rule 6). Replace both with your own facts.

| Crossing | Seen in corpus | Public example                                 | Control at the crossing                                                                | Who should see it                                                         | Authority needed to test                                                              | Jurisdiction on each side                       |
|----------|----------------|------------------------------------------------|----------------------------------------------------------------------------------------|---------------------------------------------------------------------------|---------------------------------------------------------------------------------------|-------------------------------------------------|
| H / I    | 7              | Twitter, 2020 <sup>38</sup>                    | Identity proofing at the service desk before any reset or MFA change                   | Service desk lead and identity team, from joined ticket and sign-in logs  | Your own, or the outsourcer's written participation if the desk is outsourced         | Yours; the outsourcer's if the desk sits abroad |
| H / N    | 11             | Sony Pictures, 2014 <sup>264</sup>             | Attachments and links opened in isolation; application allow-listing on workstations   | SOC, from mail gateway and endpoint telemetry                             | Your own, with HR and legal sign-off for testing staff                                | Yours                                           |
| I / N    | 30             | Change Healthcare, 2024 <sup>17</sup>          | Phishing-resistant MFA on every remote access path                                     | SOC, from VPN and gateway sign-ins matched to device identity             | Your own                                                                              | Yours                                           |
| N / I    | 22             | Nevada state government, 2025 <sup>118</sup>   | Tiered administration and credential isolation on servers and workstations             | Identity team and SOC, from directory and endpoint telemetry              | Your own                                                                              | Yours                                           |
| E / N    | 15             | Rackspace hosted Exchange, 2022 <sup>265</sup> | Segmentation behind internet-facing systems; patch windows measured in days            | SOC, from logs shipped off the appliance                                  | Your own, and the host's if the system is hosted                                      | Yours; the host's if hosted abroad              |
| I / C    | 15             | Microsoft corporate email, 2024 <sup>266</sup> | MFA and conditional access on every tenant, including test tenants                     | Cloud security, from tenant sign-in and application consent logs          | Your own, within the cloud provider's testing policy                                  | Yours; the cloud region's                       |
| E / I    | 8              | MITRE, 2024 <sup>22</sup>                      | No privileged directory credentials stored on appliances                               | Identity team, from service-account sign-ins that originate on appliances | Your own                                                                              | Yours                                           |
| N / V    | 7              | Bangladesh Bank, 2016 <sup>267</sup>           | Payment systems isolated from office IT, with independent confirmation of instructions | Payments operations and SOC, from payment system and network logs         | Your own; the payment operator's for anything past your side, and never live payments | Yours; the payment network's                    |

| Crossing                        | Seen in corpus | Public example                                                                       | Control at the crossing                                                                                     | Who should see it                                                                  | Authority needed to test                                                                               | Jurisdiction on each side                       |
|---------------------------------|----------------|--------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| I / V                           | 2              | RAKBANK, 2013, at its card processor <sup>145</sup>                                  | Separate credentials for card and payment administration, with limits changed only under dual control       | Payments operations, from the processor's administration logs shared with the bank | The processor's written participation                                                                  | Yours; the processor's                          |
| I <sub>s</sub>   / N            | 21             | Target, 2013 <sup>8</sup>                                                            | Supplier accounts limited to named systems, with MFA and expiry                                             | Supplier owner and SOC, from supplier-account activity                             | The supplier's written participation, and its regulator's if it is licensed                            | Yours; the supplier's                           |
| C <sub>s</sub>   / N            |                | US Treasury, 2024 <sup>75</sup>                                                      | Remote-support tools with per-session approval by the customer                                              | SOC, from remote-support session logs held by the customer                         | The supplier's written participation                                                                   | Yours; the supplier's                           |
| C <sub>s</sub>   C              |                | Salesloft Drift, 2025 <sup>77</sup>                                                  | Least-privilege scopes for integration tokens, with inventory and rotation                                  | SaaS owner, from integration API activity                                          | Both providers' participation, within their testing policies                                           | Each provider's; yours                          |
| I <sub>p</sub>   I              | 4              | Okta, 2023 <sup>268</sup>                                                            | Corporate credentials barred from personal browser profiles and devices                                     | Identity team, from sign-ins by unmanaged devices                                  | The employee's consent; personal devices are otherwise out of bounds                                   | The employee's; yours                           |
| P / N                           | 2              | APT28 "nearest neighbor", 2024 <sup>100</sup>                                        | Wireless and port access held to the same authentication as remote access                                   | Facilities and network team, from wireless and switch-port logs                    | Letters from the site owner and the tenant                                                             | The site's                                      |
| N / O                           | 1              | Triton, 2017 <sup>110</sup>                                                          | Engineering workstations isolated; safety systems in locked program mode                                    | OT security and plant operations, from controller change logs                      | The plant owner's; replica or lab only                                                                 | The plant's                                     |
| I <sub>t</sub>   / H            |                | North Korean IT workers, 2023 <sup>50</sup>                                          | In-person identity checks at hiring and device shipping to verified addresses                               | HR, security and IT, from onboarding and device enrollment records                 | HR and legal sign-off, and any staffing agency's participation                                         | Yours; the worker's claimed and actual location |
| A / C                           | 0              | EchoLeak, CVE-2025-32711, 2025, a flaw rather than an intrusion <sup>183</sup>       | Assistants that read outside content cannot send data out or act without approval                           | SaaS owner and SOC, from assistant and tenant audit logs                           | Your own, within the AI provider's testing terms                                                       | Yours; the AI provider's                        |
| P <sub>s</sub> / C <sub>s</sub> | 0              | AWS Bahrain and UAE regions, 2026, harm rather than an intrusion step <sup>121</sup> | Backups restorable in another region or provider, with the destination agreed with the regulator in advance | Cloud owner and continuity lead, from restore test records                         | Not testable on live facilities; test the restore, with the regulator's agreement on where data may go | The region's; where you restore                 |

Rows with a blank count have none of their own. The two supplier rows share the 21 supplier-to-victim changes counted on the I<sub>s</sub> | / N row, and I<sub>t</sub> | / H changes domain and owner at once, which neither count separates. Rows with a

count of 0 do not occur in the corpus. A / C is a published flaw, and P<sub>s</sub> / C<sub>s</sub> records harm that reached a supplier's cloud through its buildings. Their examples show that the crossings exist.

# Blank register

[illegible]

## Appendix D

# The incident corpus

All 78 cases, in order of year, with each path in bend notation and the sources it was coded from.

Each case file in the published data holds the full coding, with every step's domain, owner, action and sources, plus the actor, objective, discoverer, dwell time, impact and the coder's notes. Confidence is high when primary sources

cover every step, medium when press fills some gaps, and low when the account rests mostly on press. Year is usually the year of disclosure; the case files also record first access where the sources give it.

| Year | Case                                                                                         | Sector and place                 | Path                                                              | Confidence |
|------|----------------------------------------------------------------------------------------------|----------------------------------|-------------------------------------------------------------------|------------|
| 2013 | Bank Muscat prepaid travel card cash-out via card processor <sup>16,269,270</sup>            | Finance, Oman                    | ? <sub>s</sub> / I <sub>s</sub> / V <sub>s</sub>   V <sub>t</sub> | Medium     |
| 2013 | Carbanak campaign against banks <sup>271</sup>                                               | Finance, Global                  | H / N . N / V                                                     | Medium     |
| 2013 | RAKBANK prepaid card cash-out via card processor <sup>15,145</sup>                           | Finance, United Arab Emirates    | ? <sub>s</sub> / I <sub>s</sub> / V <sub>s</sub>   V <sub>t</sub> | Medium     |
| 2013 | Target <sup>8,9</sup>                                                                        | Retail, US                       | H <sub>s</sub>   / I / N / V / N                                  | Medium     |
| 2014 | Anthem <sup>272</sup>                                                                        | Healthcare, US                   | H / N . N                                                         | Medium     |
| 2014 | Marriott International (Starwood guest reservation database) <sup>14</sup>                   | Hospitality, Global              | E / N / I   I <sub>s</sub>   / N                                  | High       |
| 2014 | Sony Pictures Entertainment <sup>264,273</sup>                                               | Media, US                        | H / N . N                                                         | High       |
| 2014 | The Home Depot <sup>67,274</sup>                                                             | Retail, US                       | ? <sub>s</sub>   / I / N / V                                      | High       |
| 2014 | US Office of Personnel Management <sup>68,275</sup>                                          | Government, US                   | ? <sub>s</sub>   / I / N   N <sub>s</sub>                         | Medium     |
| 2015 | Bangladesh Bank <sup>267,276</sup>                                                           | Finance, Asia                    | H / N / V                                                         | High       |
| 2015 | Ukrainian regional power distribution companies <sup>108,277</sup>                           | Energy, Europe                   | H / N / I / O                                                     | Medium     |
| 2016 | Canadian Centre for Ethics in Sport (GRU close-access operation) <sup>98</sup>               | Other, Other                     | P <sub>t</sub>   / N / I / N                                      | Medium     |
| 2016 | Qatar National Bank data leak <sup>278,279,144</sup>                                         | Finance, Qatar                   | E / N                                                             | Low        |
| 2016 | Shamoon 2 wiper attacks on Saudi organizations <sup>280,281,282</sup>                        | Government, Saudi Arabia         | H / N / I / N                                                     | Medium     |
| 2016 | Ukrainian power grid (Industroyer) <sup>109,283</sup>                                        | Energy, Europe                   | ? <sub>u</sub>   / N / O                                          | Medium     |
| 2017 | DarkVishnya attacks on Eastern European banks <sup>99</sup>                                  | Finance, Europe                  | H / P / I / N / V                                                 | Medium     |
| 2017 | Equifax <sup>284,285</sup>                                                                   | Finance, US                      | E / I / N . N                                                     | High       |
| 2017 | NotPetya via M.E.Doc (organizations using M.E.Doc) <sup>70,286,287</sup>                     | Other, Global                    | I <sub>s</sub> / N <sub>s</sub>   N / I / N                       | High       |
| 2017 | Qatar News Agency website compromise <sup>288,142</sup>                                      | Media, Qatar                     | E / N / I / E                                                     | Medium     |
| 2017 | Triton safety-system attack on a Saudi petrochemical facility <sup>110,112,113,111,114</sup> | Energy, Saudi Arabia             | E / N / I / O                                                     | High       |
| 2018 | DNS hijacking of UAE government domains (DNSpionage) <sup>289,290,291,292</sup>              | Government, United Arab Emirates | I <sub>u</sub>   / E   / I <sub>t</sub>   I                       | Medium     |
| 2019 | Bapco (Dustman wiper) <sup>293,294</sup>                                                     | Energy, Bahrain                  | E / I / N                                                         | Medium     |
| 2019 | Capital One <sup>31,295,296,297</sup>                                                        | Finance, US                      | E / C                                                             | High       |
| 2019 | Norsk Hydro <sup>298,11</sup>                                                                | Manufacturing, Europe            | ? <sub>c</sub>   / H / N / I / N                                  | Medium     |

| Year | Case                                                                                                           | Sector and place                 | Path                                                     | Confidence |
|------|----------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------------------------------|------------|
| 2019 | Tortoiseshell intrusions at Saudi IT service providers <sup>299</sup>                                          | Technology, Saudi Arabia         | E / N / I                                                | Low        |
| 2020 | SolarWinds Orion supply chain compromise (generic documented victim path) <sup>71</sup>                        | Government, US                   | ? <sub>s</sub> / N <sub>s</sub>   N / I / C              | Medium     |
| 2020 | Twitter <sup>38</sup>                                                                                          | Technology, US                   | H / I / N / H / N   / I <sub>c</sub>   / H <sub>t</sub>  | High       |
| 2020 | xHunt campaign against Kuwaiti organizations <sup>300,301</sup>                                                | Government, Kuwait               | ? / E / N                                                | Medium     |
| 2021 | Colonial Pipeline <sup>10</sup>                                                                                | Energy, US                       | I / N                                                    | High       |
| 2021 | Kaseya VSA managed service providers and their customers <sup>74,302</sup>                                     | Technology, Global               | E <sub>s</sub>   / N                                     | Medium     |
| 2021 | Microsoft Exchange Server zero-day exploitation by HAFNIUM (generic documented victim path) <sup>303</sup>     | Other, US                        | E / N                                                    | Medium     |
| 2021 | Saudi Aramco data held by contractors <sup>143</sup>                                                           | Energy, Saudi Arabia             | ? <sub>s</sub>   / H                                     | Low        |
| 2022 | Cisco <sup>304</sup>                                                                                           | Technology, US                   | I <sub>p</sub>   / H / I / N / C                         | High       |
| 2022 | LastPass <sup>305,306</sup>                                                                                    | Technology, US                   | N / I / C   / N <sub>u</sub>   / I / C                   | Medium     |
| 2022 | Medibank <sup>307</sup>                                                                                        | Finance, Other                   | N <sub>p</sub>   / I / N                                 | High       |
| 2022 | Okta (via support contractor Site1) <sup>69,308</sup>                                                          | Technology, US                   | ? <sub>s</sub> / N <sub>s</sub>   / I   / C <sub>c</sub> | Medium     |
| 2022 | Rackspace Hosted Exchange <sup>265,309,310</sup>                                                               | Technology, US                   | E / N                                                    | High       |
| 2022 | Uber <sup>311</sup>                                                                                            | Transport, US                    | N <sub>p</sub>   / I / C                                 | High       |
| 2022 | Viasat KA-SAT network <sup>312,117</sup>                                                                       | Telecom, Europe                  | E <sub>s</sub> / N <sub>s</sub>   N <sub>c</sub>         | High       |
| 2023 | 3CX (via Trading Technologies X_TRADER) <sup>72</sup>                                                          | Technology, Global               | N <sub>s</sub>   N / I / N   N <sub>c</sub>              | High       |
| 2023 | Boeing Distribution Inc. <sup>25</sup>                                                                         | Manufacturing, US                | E / I / N                                                | Medium     |
| 2023 | Caesars Entertainment <sup>32</sup>                                                                            | Hospitality, US                  | H <sub>s</sub>   / N                                     | Medium     |
| 2023 | Kuwait Ministry of Finance <sup>313</sup>                                                                      | Government, Kuwait               | ? / N / H                                                | Low        |
| 2023 | Microsoft Exchange Online mailboxes of US government and other organizations (Storm-0558) <sup>79</sup>        | Government, Global               | ? <sub>s</sub> / I <sub>s</sub>   / C                    | High       |
| 2023 | Okta customer support case management system <sup>268,314</sup>                                                | Technology, US                   | I <sub>p</sub>   I / C   / I <sub>c</sub>                | High       |
| 2023 | US companies employing North Korean IT workers via Christina Chapman laptop farm <sup>50</sup>                 | Technology, US                   | I <sub>t</sub>   / H / P / V                             | High       |
| 2023 | US water and wastewater facilities with Unitronics PLCs (CyberAv3ngers) <sup>115</sup>                         | Utilities, US                    | I / O                                                    | Medium     |
| 2024 | Change Healthcare (UnitedHealth Group) <sup>17</sup>                                                           | Healthcare, US                   | I / N                                                    | High       |
| 2024 | Earth Simnavaz (APT34) intrusions into Gulf government bodies <sup>315,316</sup>                               | Government, United Arab Emirates | E / N / I / N                                            | Medium     |
| 2024 | Kuwait Ministry of Health <sup>140,317</sup>                                                                   | Government, Kuwait               | ? / N                                                    | Low        |
| 2024 | Microsoft corporate email (Midnight Blizzard) <sup>266,18</sup>                                                | Technology, US                   | I / C                                                    | High       |
| 2024 | MITRE NERVE research network <sup>22,23</sup>                                                                  | Defense, US                      | E / I / N                                                | High       |
| 2024 | Organization A (Volexity customer) via nearby Wi-Fi networks <sup>100</sup>                                    | Other, Other                     | I   / N <sub>t</sub>   / P / N                           | High       |
| 2024 | Snowflake customer tenants (UNC5537 campaign) <sup>76</sup>                                                    | Other, Global                    | N <sub>s</sub>   / I / C                                 | Medium     |
| 2024 | US Department of the Treasury Departmental Offices (via BeyondTrust Remote Support SaaS) <sup>75,318,319</sup> | Government, US                   | C <sub>s</sub> / I <sub>s</sub>   / N                    | High       |
| 2024 | US Securities and Exchange Commission @SECGov X account <sup>48,49</sup>                                       | Government, US                   | H <sub>t</sub>   / I / C                                 | High       |
| 2024 | US telecommunications providers (Salt Typhoon campaign) <sup>320,26</sup>                                      | Telecom, US                      | E / I / N                                                | Medium     |

| Year | Case                                                                                                  | Sector and place                 | Path                                                                    | Confidence |
|------|-------------------------------------------------------------------------------------------------------|----------------------------------|-------------------------------------------------------------------------|------------|
| 2025 | AI-orchestrated espionage campaign using Claude Code (GTG-1002) <sup>164</sup>                        | Other, Global                    | A <sub>t</sub>   / E / I / N                                            | Medium     |
| 2025 | Bybit cold wallet theft via Safe(Wallet) <sup>1,3,2,4</sup>                                           | Finance, Global                  | H <sub>s</sub> / I <sub>s</sub> / C <sub>s</sub>   / V                  | Medium     |
| 2025 | Coinbase bribed support agents <sup>42,321</sup>                                                      | Finance, US                      | H <sub>s</sub>   / N   / H <sub>c</sub>   H                             | High       |
| 2025 | Collins Aerospace (RTX) MUSE airport passenger processing <sup>12,13,322</sup>                        | Transport, Europe                | ? / N                                                                   | Low        |
| 2025 | F5 BIG-IP source code and vulnerability data theft <sup>323,324</sup>                                 | Technology, US                   | ? / N                                                                   | Medium     |
| 2025 | Marks & Spencer <sup>5,6,325</sup>                                                                    | Retail, UK                       | H <sub>s</sub>   / I / N                                                | Medium     |
| 2025 | npm maintainer phishing and poisoned packages (chalk, debug and others) <sup>326,327</sup>            | Technology, Global               | H <sub>s</sub> / I <sub>s</sub> / N <sub>s</sub>   N   / V <sub>c</sub> | Medium     |
| 2025 | Office of the Comptroller of the Currency email compromise <sup>328</sup>                             | Government, US                   | ? / I / C                                                               | Medium     |
| 2025 | On-premises SharePoint 'ToolShell' exploitation leading to ransomware (Storm-2603) <sup>329,330</sup> | Other, Global                    | E / N / I / N                                                           | Medium     |
| 2025 | Oracle E-Business Suite zero-day data theft and CL0P-branded extortion <sup>331</sup>                 | Other, Global                    | E . E / H                                                               | Medium     |
| 2025 | Salesloft Drift integration and customer Salesforce tenants (UNC6395) <sup>77,332,333</sup>           | Technology, Global               | I <sub>s</sub> / C <sub>s</sub>   / I / C                               | High       |
| 2025 | State of Nevada statewide ransomware incident <sup>118</sup>                                          | Government, US                   | H / N / I / N                                                           | High       |
| 2026 | 8x8 (Salesforce data via Klue integration) <sup>78</sup>                                              | Technology, US                   | C <sub>s</sub>   C                                                      | Medium     |
| 2026 | Bitget exchange hot wallets <sup>334,335</sup>                                                        | Finance, Other                   | ? <sub>u</sub>   / N / V                                                | Low        |
| 2026 | European Commission Europa web hosting cloud account (Trivy supply chain) <sup>73,81,80</sup>         | Government, Europe               | N <sub>s</sub> / I <sub>s</sub> / N <sub>s</sub>   N / C . C            | High       |
| 2026 | Evertec (Puerto Rico bank client data via support platform) <sup>336</sup>                            | Finance, US                      | ? <sub>u</sub>   / C <sub>s</sub> . C <sub>s</sub>                      | Medium     |
| 2026 | Internet-exposed PLCs at US critical infrastructure operators <sup>19</sup>                           | Utilities, US                    | O . O . O                                                               | High       |
| 2026 | Polish combined heat and power plant <sup>337,338</sup>                                               | Energy, Europe                   | I / E / N / E / N                                                       | High       |
| 2026 | Polish wind and solar farm grid connection points <sup>116,339</sup>                                  | Energy, Europe                   | E / I / O . O                                                           | Medium     |
| 2026 | Private-sector organization supporting a UAE government sector <sup>340</sup>                         | Government, United Arab Emirates | E <sub>s</sub> / N <sub>s</sub> / H <sub>s</sub>                        | Low        |
| 2026 | Stryker <sup>341,342,343,344,345</sup>                                                                | Healthcare, US                   | ? <sub>u</sub>   / I / C                                                | Medium     |

## Appendix E

# Glossary

Terms this report defines, and the outside terms it relies on, in the sense we use them.

## Terms this report defines

|                                |                                                                                                                                                                                                                                                        |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Bend</b>                    | A change of domain or of owner between two consecutive steps of an intrusion.                                                                                                                                                                          |
| <b>Bend notation</b>           | A way of writing an intrusion as a sequence of one-letter domain codes, with a small letter for any owner other than the victim, a slash for a domain bend and a vertical bar for an owner seam. Appendix B.                                           |
| <b>Consequence bend</b>        | Harm that reached a domain the adversary never acted in, usually because the victim or someone downstream shut systems down or could not operate without them. Recorded in the corpus as impact by decision.                                           |
| <b>Domain</b>                  | The kind of system or thing an adversary acts on in one step: human, identity, edge, internal IT, cloud and SaaS, operational technology, physical, value transfer or AI systems. Chosen by asking which set of controls had to fail.                  |
| <b>Domain bend</b>             | A change of domain between two consecutive steps, written with a slash in bend notation.                                                                                                                                                               |
| <b>Owner</b>                   | Whose system, account or person a step acted on: the victim, a supplier, a customer, an employee's personal account or another party.                                                                                                                  |
| <b>Owner seam</b>              | A change of owner between two consecutive steps, written with a vertical bar in bend notation. Testing across one needs the other owner's permission.                                                                                                  |
| <b>Path</b>                    | The ordered steps of one intrusion, from first access to the objective.                                                                                                                                                                                |
| <b>Recently published flaw</b> | In our analysis of the KEV catalog, a flaw published in the year CISA listed it or the year before. Counting only these removes the effect of CISA loading older flaws in 2022.                                                                        |
| <b>Seam</b>                    | A boundary where the owner, the control set, the team that watches or the legal authority changes. A bend crosses a seam.                                                                                                                              |
| <b>Seam coverage</b>           | The share of registered crossings exercised end to end in a test during the last twelve months.                                                                                                                                                        |
| <b>Seam register</b>           | A list of crossings, one row each, with a named owner on each side, the control at the crossing, who would see it, what stops beyond it, when it was last tested, whose permission a test needs and the jurisdiction on each side. Appendix C.         |
| <b>Seam-led testing</b>        | Planning tests from the crossings an adversary would need rather than from an inventory of systems. Chapter 12.                                                                                                                                        |
| <b>Single-domain test</b>      | A test scoped to one domain, or one domain and identity, such as an external or internal penetration test, a cloud configuration review, a phishing exercise, a physical test or an OT assessment. Appendix A lists the domains each is credited with. |
| <b>Telemetry seam</b>          | The boundary beyond which a provider cannot see how its product is used, for example an AI developer that sees misuse of its hosted models but not of models run elsewhere.                                                                            |
| <b>Typical red team scope</b>  | In our coding, the victim's own people and systems in human, identity, edge, internal IT, cloud and SaaS, and value transfer. Appendix A.                                                                                                              |
| <b>Walkable</b>                | A path is walkable by a test if every known step falls within the test's domains and none is on another owner's systems or people.                                                                                                                     |

## Outside terms

|                           |                                                                                                                                                                                                                                                                    |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Assumed-compromise test   | A test that starts from access an adversary would already hold, such as a working session on a VPN, instead of first trying to break in.                                                                                                                           |
| Availability zone         | One or more data centers within a cloud provider's region, with separate power and networking. A region has several.                                                                                                                                               |
| BEC                       | Business email compromise: fraud in which a person is deceived, usually by email, into sending money or changing payment details.                                                                                                                                  |
| Bilateral recognition     | An agreement between two supervisors, often the home and host supervisors of a banking group, that a test run under one counts for the other.                                                                                                                      |
| BOD 26-04                 | CISA Binding Operational Directive 26-04 of June 2026. For KEV entries marked for forensic triage, it requires US federal civilian agencies to remediate or mitigate within three days and to carry out a forensic triage of the asset to check it for compromise. |
| CBB                       | The Central Bank of Bahrain.                                                                                                                                                                                                                                       |
| CBEST                     | The Bank of England's framework for intelligence-led testing of UK financial firms.                                                                                                                                                                                |
| CBK                       | The Central Bank of Kuwait.                                                                                                                                                                                                                                        |
| CBO                       | The Central Bank of Oman.                                                                                                                                                                                                                                          |
| CBUAE                     | The Central Bank of the United Arab Emirates.                                                                                                                                                                                                                      |
| CIRCIA                    | The US Cyber Incident Reporting for Critical Infrastructure Act of 2022. Its reporting duties apply only once CISA's final rule takes effect.                                                                                                                      |
| Close access              | An operation in which the adversary is physically near the target, for example within radio range of its Wi-Fi.                                                                                                                                                    |
| Cohen's kappa             | A measure of agreement between two coders that corrects for the agreement expected by chance. One is perfect agreement, zero is chance.                                                                                                                            |
| CORF                      | The Central Bank of Kuwait's cyber and operational resilience framework, issued December 3, 2025.                                                                                                                                                                  |
| CVE                       | Common Vulnerabilities and Exposures, the public identifier assigned to a disclosed vulnerability.                                                                                                                                                                 |
| DFSA                      | The Dubai Financial Services Authority, the regulator of firms in the Dubai International Financial Centre.                                                                                                                                                        |
| DMARC                     | A DNS record in which a domain owner tells mail receivers what to do with mail that fails authentication: nothing (p=none), quarantine it or reject it.                                                                                                            |
| DNSSEC                    | Signatures on DNS records that let resolvers check an answer came from the domain's owner, valid only when the parent zone holds a matching DS record.                                                                                                             |
| DORA                      | The EU Digital Operational Resilience Act, Regulation (EU) 2022/2554, which includes threat-led penetration testing for certain financial entities and pooled tests of shared ICT providers.                                                                       |
| EPSS                      | The Exploit Prediction Scoring System, run by FIRST, which estimates the probability that a published vulnerability will be exploited in the next 30 days.                                                                                                         |
| FEER                      | SAMA's Financial Entities Ethical Red Teaming framework for Saudi financial institutions.                                                                                                                                                                          |
| GNSS interference         | Jamming or spoofing of satellite navigation signals such as GPS.                                                                                                                                                                                                   |
| Home and host supervisors | For a banking group, the supervisor of the parent (home) and the supervisor of an entity in another country (host).                                                                                                                                                |
| Item 1.05                 | The Form 8-K item under which US public companies disclose material cybersecurity incidents, within four business days of determining materiality.                                                                                                                 |
| Item 8.01                 | The "Other Events" item of Form 8-K. After an SEC staff statement of May 21, 2024, companies could disclose incidents not yet determined to be material under it rather than under Item 1.05.                                                                      |

|                                       |                                                                                                                                                                                                    |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Jackpotting</b>                    | Forcing a cash machine to dispense money, usually by opening it and tampering with its hard drive.                                                                                                 |
| <b>KEV</b>                            | CISA's Known Exploited Vulnerabilities catalog: vulnerabilities with reliable evidence of exploitation, an assigned CVE and a clear remediation.                                                   |
| <b>Laptop farm</b>                    | A residence in one country where company laptops are kept and connected so that workers elsewhere can use them while appearing local.                                                              |
| <b>MTA-STS</b>                        | A policy by which a domain tells sending mail servers to require authenticated, encrypted delivery.                                                                                                |
| <b>NCA</b>                            | Saudi Arabia's National Cybersecurity Authority.                                                                                                                                                   |
| <b>NYDFS</b>                          | The New York State Department of Financial Services, which supervises banks and insurers licensed in New York.                                                                                     |
| <b>Open-weight model</b>              | An AI model whose trained parameters are published, so that anyone can run and modify it on their own hardware.                                                                                    |
| <b>Passkey</b>                        | A sign-in credential based on a key pair held on a device or in a password manager, which a fake site cannot capture.                                                                              |
| <b>Pooled test</b>                    | One threat-led test of a shared supplier run on behalf of several of its customers, as DORA allows for ICT providers.                                                                              |
| <b>Prompt injection</b>               | Content that an AI system reads as data but treats as instructions, causing it to act or disclose information against its operator's intent.                                                       |
| <b>Purple teaming or replay</b>       | A session after a red team test in which attackers and defenders walk through the timeline together.                                                                                               |
| <b>QCB</b>                            | Qatar Central Bank.                                                                                                                                                                                |
| <b>Red team</b>                       | An objective-led test in which the testers act as a real adversary would, across whatever systems and people the rules of engagement allow.                                                        |
| <b>SAMA</b>                           | The Saudi Central Bank, formerly the Saudi Arabian Monetary Authority.                                                                                                                             |
| <b>security.txt</b>                   | A file at a fixed web address, defined in RFC 9116, that tells outsiders how to report a security problem. It must include a contact and an expiry date.                                           |
| <b>SIM swap</b>                       | Moving a victim's phone number to a SIM card the adversary controls, usually through a carrier's store or support staff.                                                                           |
| <b>SPF</b>                            | A DNS record listing the servers allowed to send mail for a domain. A record ending in -all tells receivers to fail all other senders.                                                             |
| <b>Threat-led penetration testing</b> | Testing in which intelligence on real adversaries shapes the scenarios, as in CBEST, TIBER-EU and DORA.                                                                                            |
| <b>TIBER-EU</b>                       | The European Central Bank's framework for threat intelligence-based ethical red teaming.                                                                                                           |
| <b>TLP</b>                            | The Traffic Light Protocol, markings that tell recipients how widely they may share information. TLP:AMBER limits sharing to the recipient's organization and its clients on a need-to-know basis. |
| <b>Uncensored model</b>               | An AI model, or a service built on one, offered without the safety restrictions of mainstream providers, often a wrapper around a commercial model or a modified open-weight model.                |
| <b>Vishing</b>                        | Voice phishing: deceiving someone by telephone.                                                                                                                                                    |
| <b>Wilson interval</b>                | A range for a share computed from a count. We give 95 percent Wilson intervals as a scale for how much a count could move by chance, not as estimates for all intrusions.                          |
| <b>Zero-day</b>                       | A flaw exploited before its vendor has disclosed or fixed it.                                                                                                                                      |

# Notes

Sources are numbered in the order they are first cited, with the date each was accessed.

- 1 North Korea Responsible for \$1.5 Billion Bybit Hack (Alert I-022625-PSA). FBI Internet Crime Complaint Center, February 26, 2025. <https://www.ic3.gov/PSA/2025/PSA250226> Accessed September 28, 2026.
- 2 Bybit – What We Know So Far. Sygnia, March 16, 2025. <https://www.sygnia.co/blog/sygnia-investigation-bybit-hack/> Accessed September 28, 2026.
- 3 Safe{Wallet} statement and Mandiant preliminary findings on the Bybit incident. Safe{Wallet} (Safe Ecosystem Foundation), March 6, 2025. <https://x.com/safe/status/18976635149756499> Accessed September 28, 2026.
- 4 Lazarus hacked Bybit via breached Safe{Wallet} developer machine. BleepingComputer, February 26, 2025. Press report. <https://www.bleepingcomputer.com/news/security/lazarus-hacked-bybit-via-breached-safe-wallet-developer-machine/> Accessed September 28, 2026.
- 5 Business and Trade Sub-Committee on Economic Security, Arms and Export Controls: oral evidence (Archie Norman, M&S). UK Parliament, July 8, 2025. <https://committees.parliament.uk/oralevidence/16268/html/> Accessed September 28, 2026.
- 6 M&S confirms social engineering led to massive ransomware attack. BleepingComputer, July 8, 2025. Press report. <https://www.bleepingcomputer.com/news/security/mands-confirms-social-engineering-led-to-massive-ransomware-attack/> Accessed September 28, 2026.
- 7 MITRE ATT&CK. The MITRE Corporation, September 28, 2026. <https://attack.mitre.org/> Accessed September 28, 2026.
- 8 A "Kill Chain" Analysis of the 2013 Target Data Breach, Majority Staff Report for Chairman Rockefeller, printed in S. Hrg. 113-531. U.S. Senate Committee on Commerce, Science, and Transportation (via U.S. Government Publishing Office), March 26, 2014. <https://www.govinfo.gov/content/pkg/CHRG-113shrg92594/html/CHRG-113shrg92594.htm> Accessed September 28, 2026.
- 9 Testimony of John J. Mulligan, Executive Vice President and CFO, Target Corporation, in S. Hrg. 113-531, Protecting Personal Consumer Information from Cyber Attacks and Data Breaches. U.S. Senate Committee on Commerce, Science, and Transportation (via U.S. Government Publishing Office), March 26, 2014. <https://www.govinfo.gov/content/pkg/CHRG-113shrg92594/html/CHRG-113shrg92594.htm> Accessed September 28, 2026.
- 10 Testimony of Joseph Blount, President and CEO, Colonial Pipeline Company, before the House Committee on Homeland Security. US House of Representatives, Committee on Homeland Security, June 9, 2021. <https://docs.house.gov/meetings/HM/HM00/20210609/112689/HHRG-117-HM00-Wstate-BlountJ-20210609.pdf> Accessed September 28, 2026.
- 11 LockerGoga Ransomware Attack on Norsk Hydro 2019: Precursor Analysis Report. Idaho National Laboratory, CyOTE program, US Department of Energy, September 30, 2022. [https://cyote.inl.gov/content/uploads/24/2025/12/CyOTE-Case-Study\\_Norsk-Hydro.pdf](https://cyote.inl.gov/content/uploads/24/2025/12/CyOTE-Case-Study_Norsk-Hydro.pdf) Accessed September 28, 2026.
- 12 RTX Corporation Form 8-K, Item 8.01 Other Events. U.S. Securities and Exchange Commission (RTX filing), September 19, 2025. <https://www.sec.gov/Archives/edgar/data/101829/000010182925000036/rtx-20250919.htm> Accessed September 28, 2026.
- 13 UK arrest following aerospace cyber incident. UK National Crime Agency, September 24, 2025. <https://www.nationalcrimeagency.gov.uk/news/uk-arrest-following-aerospace-cyber-incident> Accessed September 28, 2026.
- 14 Penalty Notice: Marriott International Inc. UK Information Commissioner's Office, October 30, 2020. <https://ico.org.uk/media2/migrated/2618524/marriott-international-inc-mpn-20201030.pdf> Accessed September 28, 2026.
- 15 Eight Members Of New York Cell Of Cybercrime Organization Indicted In \$45 Million Cybercrime Campaign. U.S. Attorney's Office, Eastern District of New York, May 9, 2013. <https://www.justice.gov/usao-edny/pr/eight-members-new-york-cell-cybercrime-organization-indicted-45-million-cybercrime> Accessed September 28, 2026.
- 16 Eight Members Of New York Cell Of Cybercrime Organization Indicted In \$45 Million Cybercrime Campaign. U.S. Attorney's Office, Eastern District of New York, May 9, 2013. <https://www.justice.gov/usao-edny/pr/eight-members-new-york-cell-cybercrime-organization-indicted-45-million-cybercrime> Accessed September 28, 2026.
- 17 Testimony of Andrew Witty, Chief Executive Officer, UnitedHealth Group. US Senate Committee on Finance, May 1, 2024. [https://www.finance.senate.gov/imo/media/doc/0501\\_witty\\_testimony.pdf](https://www.finance.senate.gov/imo/media/doc/0501_witty_testimony.pdf) Accessed September 28, 2026.
- 18 Midnight Blizzard: Guidance for responders on nation-state attack. Microsoft Threat Intelligence, January 25, 2024. <https://www.microsoft.com/en-us/security/blog/2024/01/25/midnight-blizzard-guidance-for-responders-on-nation-state-attack/> Accessed September 28, 2026.
- 19 Iranian-Affiliated Cyber Actors Exploit Programmable Logic Controllers Across US Critical Infrastructure (AA26-097A, updated 22 July 2026). CISA, FBI and partners, July 22, 2026. <https://www.cisa.gov/news-events/cybersecurity-advisories/aa26-097a> Accessed September 28, 2026.
- 20 Known Exploited Vulnerabilities Catalog. Cybersecurity and Infrastructure Security Agency, September 27, 2026. <https://www.cisa.gov/known-exploited-vulnerabilities-catalog> Accessed September 28, 2026.
- 21 CVE API 2.0, query hasKev. National Institute of Standards and Technology, National Vulnerability Database, September 28, 2026. <https://services.nvd.nist.gov/rest/json/cves/2.0?hasKev&resultsPerPage=2000&startIndex=0> Accessed September 28, 2026.
- 22 Technical Deep Dive: Understanding the Anatomy of a Cyber Intrusion. MITRE Center for Threat-Informed Defense (Medium), May 3, 2024. <https://medium.com/mitre-engenuity/technical-deep-dive-understanding-the-anatomy-of-a-cyber-intrusion-080bddc679f3> Accessed September 28, 2026.
- 23 Infiltrating Defenses: Abusing VMware in MITRE's Cyber Intrusion. MITRE Center for Threat-Informed Defense (Medium), May 22, 2024. <https://medium.com/mitre-engenuity/infiltrating-defenses-abusing-vmware-in-mitres-cyber-intrusion-4ea647b83f5b> Accessed September 28, 2026.

- 24 Known Exploited Vulnerabilities catalog, JSON feed, catalogVersion 2026.09.27. Cybersecurity and Infrastructure Security Agency, September 27, 2026.  
[https://www.cisa.gov/sites/default/files/feeds/known\\_exploited\\_vulnerabilities.json](https://www.cisa.gov/sites/default/files/feeds/known_exploited_vulnerabilities.json) Accessed September 28, 2026.
- 25 #StopRansomware: LockBit 3.0 Ransomware Affiliates Exploit CVE 2023-4966 Citrix Bleed Vulnerability (AA23-325A). CISA, FBI, MS-ISAC and ASD's ACSC, November 21, 2023.  
<https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-325a> Accessed September 28, 2026.
- 26 Countering Chinese State-Sponsored Actors Compromise of Networks Worldwide to Feed Global Espionage System (AA25-239A). CISA, NSA, FBI and international partners, September 3, 2025. <https://www.cisa.gov/news-events/cybersecurity-advisories/aa25-239a> Accessed September 28, 2026.
- 27 Exploit Prediction Scoring System (EPSS). FIRST. <https://www.first.org/epss/> Accessed September 28, 2026.
- 28 EPSS API, daily scores. FIRST. <https://api.first.org/data/v1/epss> Accessed September 28, 2026.
- 29 Binding Operational Directive 26-04: Prioritizing Security Updates Based on Risk. Cybersecurity and Infrastructure Security Agency, June 10, 2026. <https://www.cisa.gov/news-events/directives/bod-26-04-prioritizing-security-updates-based-risk> Accessed September 28, 2026.
- 30 Reducing the Significant Risk of Known Exploited Vulnerabilities. Cybersecurity and Infrastructure Security Agency. <https://www.cisa.gov/known-exploited-vulnerabilities-catalog/reducing-significant-risk-known-exploited-vulnerabilities> Accessed September 28, 2026.
- 31 Information on the Capital One cyber incident. Capital One. <https://www.capitalone.com/digital/facts2019/> Accessed September 28, 2026.
- 32 Caesars Entertainment, Inc. Form 8-K (Item 8.01). Caesars Entertainment / US SEC EDGAR, September 14, 2023. <https://www.sec.gov/Archives/edgar/data/1590895/000119312523235015/d537840d8k.htm> Accessed September 28, 2026.
- 33 Scattered Spider (AA23-320A), updated July 29, 2025. Cybersecurity and Infrastructure Security Agency, Federal Bureau of Investigation and partners, November 16, 2023. <https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-320a> Accessed September 28, 2026.
- 34 Post on Scattered Spider targeting the airline sector. Federal Bureau of Investigation (@FBI on X), June 27, 2025. <https://x.com/FBI/status/193874676703157456> 5 Accessed September 28, 2026.
- 35 Cyber Criminal Groups UNC6040 and UNC6395 Compromising Salesforce Instances for Data Theft and Extortion (FLASH-20250912-001). Federal Bureau of Investigation, September 12, 2025. <https://www.ic3.gov/CSA/2025/250912.pdf> Accessed September 28, 2026.
- 36 Social Engineering Tactics Targeting Healthcare and Public Health Entities and Providers. Federal Bureau of Investigation and Department of Health and Human Services, June 24, 2024. <https://www.ic3.gov/CSA/2024/240624.pdf> Accessed September 28, 2026.
- 37 Ongoing Cyberthreat Campaign Involving Targeted "Vishing" Attacks (industry letter). New York State Department of Financial Services, February 6, 2026. <https://www.dfs.ny.gov/industry-guidance/industry-letters/20260206-cybersecurity-advisory-targeted-vishing-attacks> Accessed September 28, 2026.
- 38 Twitter Investigation Report: A Report on Investigation of Twitter's July 15, 2020 Cybersecurity Incident and the Implications for Election Security. New York State Department of Financial Services, 2020-10. <https://www.dfs.ny.gov/system/files/documents/2026/07/Twitter-Investigation-Report.pdf> Accessed September 28, 2026.
- 39 M-Trends 2026: Data, Insights, and Strategies From the Frontlines. Mandiant, Google Cloud, March 24, 2026. <https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026/> Accessed September 28, 2026.
- 40 UNC6671 Rebrands: Multi-Brand Vishing Extortion Targets Financial Services and Enterprise Cloud Environments. Google Threat Intelligence Group, August 7, 2026. <https://cloud.google.com/blog/topics/threat-intelligence/unc6671-targets-financial-services-and-enterprise-cloud-environments> Accessed September 28, 2026.
- 41 United Kingdom National Charged in Connection with Multiple Cyber Attacks, Including on Critical Infrastructure. US Department of Justice, September 18, 2025. <https://www.justice.gov/opa/pr/united-kingdom-national-charged-connection-multiple-cyber-attacks-including-critical> Accessed September 28, 2026.
- 42 Coinbase Global, Inc. Form 8-K, Item 1.05 Material Cybersecurity Incident. U.S. Securities and Exchange Commission (Coinbase filing), May 14, 2025. <https://www.sec.gov/Archives/edgar/data/1679788/000167978825000094/coin-20250514.htm> Accessed September 28, 2026.
- 43 Coinbase Global, Inc. Form 10-K for fiscal year 2025. Coinbase Global, Inc. (SEC EDGAR), February 12, 2026. <https://www.sec.gov/Archives/edgar/data/1679788/000167978826000015/coin-20251231.htm> Accessed September 28, 2026.
- 44 Coinbase Q2 2025 Shareholder Letter (Form 8-K Exhibit 99.1). Coinbase Global, Inc. (SEC EDGAR), July 31, 2025. <https://www.sec.gov/Archives/edgar/data/1679788/000167978825000153/q225shareholderletter.htm> Accessed September 28, 2026.
- 45 Protecting Our Customers: Standing Up to Extortionists. Coinbase, May 15, 2025. <https://www.coinbase.com/blog/protecting-our-customers-standing-up-to-extortionists> Accessed September 28, 2026.
- 46 Review of the Attacks Associated with Lapsus\$ and Related Threat Groups. Cyber Safety Review Board, US Department of Homeland Security, July 24, 2023. [https://www.cisa.gov/sites/default/files/2023-08/CSRB\\_Lapsus\\$.508c.pdf](https://www.cisa.gov/sites/default/files/2023-08/CSRB_Lapsus$.508c.pdf) Accessed September 28, 2026.
- 47 Oregon Man Sentenced to 16 Months in Federal Prison for Abusing His Role at Mobile Phone Store to Give Customers' Info to Criminals. US Attorney's Office, Central District of California, September 8, 2026. <https://www.justice.gov/usao-cdca/pr/oregon-man-sentenced-16-months-federal-prison-abusing-his-role-mobile-phone-store-give> Accessed September 28, 2026.
- 48 Alabama Man Sentenced to 14 Months in Connection with Securities and Exchange Commission X Hack that Spiked Bitcoin Value. US Department of Justice, May 16, 2025. <https://www.justice.gov/opa/pr/alabama-man-sentenced-14-months-connection-securities-and-exchange-commission-x-hack-spiked> Accessed September 28, 2026.
- 49 SECgov X Account. US Securities and Exchange Commission, January 22, 2024. <https://www.sec.gov/secgov-x-account> Accessed September 28, 2026.
- 50 Arizona Woman Sentenced in \$17M IT Worker Fraud Scheme That Illegally Generated Revenue for North Korea. US Attorney's Office, District of Columbia (DOJ), July 24, 2025. <https://www.justice.gov/usao-dc/pr/arizona-woman-sentenced-17m-it-worker-fraud-scheme-illegally-generated-revenue-north> Accessed September 28, 2026.
- 51 Justice Department Announces Coordinated, Nationwide Actions to Combat North Korean Remote Information Technology Workers' Illicit Revenue Generation Schemes. US Department of Justice, June 30, 2025. <https://www.justice.gov/opa/pr/justice-department-announces-coordinated-nationwide-actions-combat-north-korean-remote> Accessed September 28, 2026.
- 52 Treasury Sanctions Facilitators of DPRK IT Worker Fraud Targeting U.S. Businesses (sb0416). US Department of the Treasury, March 12, 2026. <https://home.treasury.gov/news/press-releases/sb0416> Accessed September 28, 2026.

- 53 Maryland Man Sentenced for Conspiracy to Commit Wire Fraud. US Department of Justice, December 4, 2025.  
<https://www.justice.gov/opa/pr/maryland-man-sentenced-conspiracy-commit-wire-fraud> Accessed September 28, 2026.
- 54 Three sentenced for facilitating computer access in North Korean sanctions evasion scheme. US Attorney's Office, Southern District of Georgia, March 20, 2026.  
<https://www.justice.gov/usao-sdga/pr/three-sentenced-facilitating-computer-access-north-korean-sanctions-evasion-scheme> Accessed September 28, 2026.
- 55 How a North Korean Fake IT Worker Tried to Infiltrate Us. KnowBe4, July 23, 2024.  
<https://blog.knowbe4.com/how-a-north-korean-fake-it-worker-tried-to-infiltrate-us> Accessed September 28, 2026.
- 56 Internet Crime Report 2025. FBI Internet Crime Complaint Center, April 6, 2026.  
[https://www.ic3.gov/AnnualReport/Reports/2025\\_IC3Report.pdf](https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf) Accessed September 28, 2026.
- 57 Organized Financial Fraud (strategic analysis). UAE Financial Intelligence Unit, November 29, 2024.  
<https://uaefiu.gov.ae/media/1pmnbrjz/fraud-report-24-nov-final-sanitised-version-final.pdf> Accessed September 28, 2026.
- 58 Cybercrime: Lessons learned from a \$25m deepfake attack (interview with Arup CIO Rob Greig). World Economic Forum, February 4, 2025.  
<https://www.weforum.org/stories/2025/02/deepfake-ai-cybercrime-arup/> Accessed September 28, 2026.
- 59 Internet Crime Report 2023. FBI Internet Crime Complaint Center, 2024.  
[https://www.ic3.gov/AnnualReport/Reports/2023\\_IC3Report.pdf](https://www.ic3.gov/AnnualReport/Reports/2023_IC3Report.pdf) Accessed September 28, 2026.
- 60 MGM Resorts International Form 8-K. MGM Resorts International (SEC EDGAR), October 5, 2023.  
<https://www.sec.gov/Archives/edgar/data/789570/000119312523251667/d461062d8k.htm> Accessed September 28, 2026.
- 61 EDGAR full-text search and filing archives, Form 8-K and 8-K/A cybersecurity incident disclosures, December 18, 2023 to September 28, 2026. US Securities and Exchange Commission, September 28, 2026.  
<https://efts.sec.gov/LATEST/search-index> Accessed September 28, 2026.
- 62 Form 8-K, General Instructions, Item 1.05 and Instructions to Item 1.05. US Securities and Exchange Commission.  
<https://www.sec.gov/files/form8-k.pdf> Accessed September 28, 2026.
- 63 Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure: Small Entity Compliance Guide. US Securities and Exchange Commission.  
<https://www.sec.gov/resources-small-businesses/small-business-compliance-guides/cybersecurity-risk-management-strategy-governance-incident-disclosure> Accessed September 28, 2026.
- 64 Federal Home Loan Bank of New York, Form 8-K. US Securities and Exchange Commission, EDGAR, March 1, 2024.  
[https://www.sec.gov/Archives/edgar/data/1329842/000165495424002505/fhlbny\\_8k.htm](https://www.sec.gov/Archives/edgar/data/1329842/000165495424002505/fhlbny_8k.htm) Accessed September 28, 2026.
- 65 Disclosure of Cybersecurity Incidents Determined To Be Material and Other Cybersecurity Incidents. Erik Gerding, US Securities and Exchange Commission, May 21, 2024.  
<https://www.sec.gov/newsroom/speeches-statements/gerding-cybersecurity-incidents-05212024> Accessed September 28, 2026.
- 66 Cybersecurity Incident Disclosure: Form 8-K Tracker (Two-Year Update). Debevoise & Plimpton, May 21, 2026.  
<https://www.debevoisedatablog.com/2026/05/21/cybersecurity-incident-disclosure-form-8-k-tracker-two-year-update/> Accessed September 28, 2026.
- 67 The Home Depot, Inc. Form 8-K, exhibit: The Home Depot Reports Findings in Payment Data Breach Investigation. U.S. Securities and Exchange Commission (EDGAR), November 6, 2014.  
[https://www.sec.gov/Archives/edgar/data/0000354950/000035495014000042/hd\\_8kx110614.htm](https://www.sec.gov/Archives/edgar/data/0000354950/000035495014000042/hd_8kx110614.htm) Accessed September 28, 2026.
- 68 Timeline of Key Events [short version], accompanying The OPM Data Breach: How the Government Jeopardized Our National Security for More than a Generation. U.S. House Committee on Oversight and Government Reform, Majority Staff, September 7, 2016.  
<https://oversight.house.gov/wp-content/uploads/2016/09/Timeline-of-Key-Events-.pdf> Accessed September 28, 2026.
- 69 Okta Concludes its Investigation Into the January 2022 Compromise. Okta, April 20, 2022. <https://www.okta.com/blog/2022/04/okta-concludes-its-investigation-into-the-january-2022-compromise> Accessed September 28, 2026.
- 70 The MeDoc Connection. Cisco Talos, July 5, 2017. <https://blog.talosintelligence.com/the-medoc-connection/> Accessed September 28, 2026.
- 71 Advanced Persistent Threat Compromise of Government Agencies, Critical Infrastructure, and Private Sector Organizations (AA20-352A). Cybersecurity and Infrastructure Security Agency, December 17, 2020.  
<https://www.cisa.gov/news-events/cybersecurity-advisories/aa20-352a> Accessed September 28, 2026.
- 72 3CX Software Supply Chain Compromise Initiated by a Prior Software Supply Chain Compromise; Suspected North Korean Actor Responsible. Mandiant (Google Cloud), April 20, 2023.  
<https://cloud.google.com/blog/topics/threat-intelligence/3cx-software-supply-chain-compromise> Accessed September 28, 2026.
- 73 European Commission cloud breach: a supply-chain compromise. CERT-EU, April 2, 2026.  
<https://cert.europa.eu/blog/european-commission-cloud-breach-trivy-supply-chain> Accessed September 28, 2026.
- 74 Kaseya Ransomware Attack: Guidance for Affected MSPs and their Customers. Cybersecurity and Infrastructure Security Agency, 2021-07. <https://www.cisa.gov/news-events/news/kaseya-ransomware-attack-guidance-affected-msps-and-their-customers> Accessed September 28, 2026.
- 75 Letter from Aditi Hardikar, Assistant Secretary for Management, to Chairman Brown and Ranking Member Scott. US Department of the Treasury (copy hosted by Nextgov), December 30, 2024.  
[https://www.nextgov.com/media/general/2024/12/12.30.2024\\_letter\\_to\\_chairman\\_brown\\_and\\_ranking\\_member\\_scott.pdf](https://www.nextgov.com/media/general/2024/12/12.30.2024_letter_to_chairman_brown_and_ranking_member_scott.pdf) Accessed September 28, 2026.
- 76 UNC5537 Targets Snowflake Customer Instances for Data Theft and Extortion. Mandiant (Google Cloud), June 10, 2024.  
<https://cloud.google.com/blog/topics/threat-intelligence/unc5537-snowflake-data-theft-extortion> Accessed September 28, 2026.
- 77 Widespread Data Theft Targets Salesforce Instances via Salesloft Drift. Google Threat Intelligence Group, August 26, 2025.  
<https://cloud.google.com/blog/topics/threat-intelligence/data-theft-salesforce-instances-via-salesloft-drift> Accessed September 28, 2026.
- 78 8x8, Inc. Form 8-K (Item 1.05). 8x8 / SEC EDGAR, June 17, 2026.  
<https://www.sec.gov/Archives/edgar/data/0001023731/000102373126000084/eght-20260617.htm> Accessed September 28, 2026.
- 79 Review of the Summer 2023 Microsoft Exchange Online Intrusion. Cyber Safety Review Board (DHS/CISA), March 20, 2024.  
<https://www.cisa.gov/sites/default/files/2025-03/CSRBReviewOfTheSummer2023MEOIntrusion508.pdf> Accessed September 28, 2026.
- 80 Trivy Supply Chain Attack: What You Need to Know. Aqua Security, March 22, 2026.  
<https://www.aquasec.com/blog/trivy-supply-chain-attack-what-you-need-to-know/> Accessed September 28, 2026.
- 81 Commission responds to cyber-attack on its Europa web platform (IP/26/748). European Commission, March 27, 2026.  
[https://ec.europa.eu/commission/presscorner/detail/en/ip\\_26\\_748](https://ec.europa.eu/commission/presscorner/detail/en/ip_26_748) Accessed September 28, 2026.

- 82 AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise. Microsoft Threat Intelligence, October 21, 2026.  
<https://www.microsoft.com/en-us/security/blog/2026/10/21/storm-0940-uses-artex-agentic-ai-connectwise-supply-chain-compromise/> Accessed October 26, 2026.
- 83 Thematic findings from the 2024 Cyber Stress Test. Bank of England, July 9, 2025.  
<https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/letter/2024/thematic-findings-cyber.pdf> Accessed October 26, 2026.
- 84 Cloudflare deploys protections against php2shell vulnerability. Cloudflare, October 21, 2026. <https://blog.cloudflare.com/php-vulnerabilities-waf-protection/> Accessed October 26, 2026.
- 85 Update: PHP vulnerabilities remain a reason for immediate patching. National Cyber Security Centre, Netherlands, October 23, 2026.  
<https://www.ncsc.nl/en/alerts/update-php-vulnerabilities-remain-a-reason-for-immediate-patching> Accessed October 26, 2026.
- 86 Cyber Security: PHP Vulnerability, House of Lords, volume 860. UK Parliament, Hansard, October 22, 2026.  
<https://hansard.parliament.uk/Lords/2026-10-22/debates/14CB4C6D-4E16-4377-8FCF-C2CF77BB5CC0/details> Accessed October 26, 2026.
- 87 Security advisory: Critical remote code execution vulnerabilities, update now. The PHP Group, php.net, October 21, 2026.  
<https://www.php.net/archive/2026.php#2026-10-21-5> Accessed October 26, 2026.
- 88 Today's PHP Security Release, and How We're Strengthening PHP Security. The PHP Foundation, October 21, 2026.  
<https://thephp.foundation/blog/2026/10/21/todays-php-security-release-and-how-were-strengthening-php-security/> Accessed October 26, 2026.
- 89 Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA). Official Journal of the European Union, December 14, 2022. <https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng> Accessed September 28, 2026.
- 90 Commission Delegated Regulation (EU) 2025/1190 (RTS on threat-led penetration testing). Official Journal of the European Union, June 18, 2025. [https://eur-lex.europa.eu/eli/reg\\_del/2025/1190/oj/eng](https://eur-lex.europa.eu/eli/reg_del/2025/1190/oj/eng) Accessed September 28, 2026.
- 91 CBB Rulebook Volume 1, OM-5.5.57 to OM-5.5.60. Central Bank of Bahrain, 2022-04.  
<https://cbben.thomsonreuters.com/rulebook/om-5557> Accessed September 28, 2026.
- 92 Technology Risks (QCB cyber security circular for banks). Qatar Central Bank, 2018-01.  
<https://www.qcb.gov.qa/Documents/InformationSecurity/Technology%20Risks%20Regulation%20Banks.pdf> Accessed September 28, 2026.
- 93 Cyber and Operational Resilience Framework for All Local Banks and Financial Institutions (CORF), Version 1.0. Central Bank of Kuwait, December 3, 2025.  
[https://www.cbk.gov.kw/en/images/corf-170113\\_v10\\_tcm10-170113.pdf](https://www.cbk.gov.kw/en/images/corf-170113_v10_tcm10-170113.pdf) Accessed September 28, 2026.
- 94 Circular No. 1/2026, Operational Risk Management Regulation. Central Bank of the UAE, February 3, 2026.  
[https://rulebook.centralbank.ae/sites/default/files/en\\_net\\_file\\_store/CBUAE\\_EN\\_7267\\_VER1.pdf](https://rulebook.centralbank.ae/sites/default/files/en_net_file_store/CBUAE_EN_7267_VER1.pdf) Accessed September 28, 2026.
- 95 Increase in Malware Enabled ATM Jackpotting Incidents Across United States (FLASH-20260219-001). Federal Bureau of Investigation, February 19, 2026.  
<https://www.ic3.gov/CSA/2026/260219.pdf> Accessed September 28, 2026.
- 96 Operational Technology Cybersecurity Controls (OTCC-1:2022). National Cybersecurity Authority, Saudi Arabia, 2022.  
[https://nca.gov.sa/otcc\\_en.pdf](https://nca.gov.sa/otcc_en.pdf) Accessed September 28, 2026.
- 97 NISCF Accreditation Penetration Testing Standard V1.1. National Cyber Security Agency (Qatar), February 8, 2024.  
[https://ncsa.gov.qa/shared/document\\_library/94787172324409/33000/3861/1.0~1fca5e63-9fff-4172-82c5-422b356ef221?mime\\_type=application/pdf&filename=SD.pdf](https://ncsa.gov.qa/shared/document_library/94787172324409/33000/3861/1.0~1fca5e63-9fff-4172-82c5-422b356ef221?mime_type=application/pdf&filename=SD.pdf) Accessed September 28, 2026.
- 98 U.S. Charges Russian GRU Officers with International Hacking and Related Influence and Disinformation Operations. U.S. Department of Justice, Office of Public Affairs, October 4, 2018.  
<https://www.justice.gov/archives/opa/pr/us-charges-russian-gru-officers-international-hacking-and-related-influence-and> Accessed September 28, 2026.
- 99 DarkVishnya: Banks attacked through direct connection to local network. Kaspersky Lab, Securelist (Sergey Golovanov), December 6, 2018.  
<https://securelist.com/darkvishnya/89169/> Accessed September 28, 2026.
- 100 The Nearest Neighbor Attack: How A Russian APT Weaponized Nearby Wi-Fi Networks for Covert Access. Volexity, November 22, 2024.  
<https://www.volexity.com/blog/2024/11/22/the-nearest-neighbor-attack-how-a-russian-apt-weaponized-nearby-wi-fi-networks-for-covert-access/> Accessed September 28, 2026.
- 101 GRU close access cyber operation against OPCW (MIVD press conference presentation). Netherlands Ministry of Defence, copy hosted by the US Department of Justice, October 4, 2018.  
<https://www.justice.gov/opa/page/file/1098576/download> Accessed September 28, 2026.
- 102 Camera off: Akira deploys ransomware via webcam. S-RM, March 5, 2025. <https://www.s-rminform.com/latest-thinking/camera-off-akira-deploys-ransomware-via-webcam> Accessed September 28, 2026.
- 103 Russian GRU Targeting Western Logistics Entities and Technology Companies (AA25-141A). NSA, FBI, CISA and partners, 2025-05.  
<https://www.cisa.gov/news-events/cybersecurity-advisories/aa25-141a> Accessed September 28, 2026.
- 104 Interplay between Iranian Targeting of IP Cameras and Physical Warfare in the Middle East. Check Point Research, March 4, 2026.  
<https://research.checkpoint.com/2026/interplay-between-iranian-targeting-of-ip-cameras-and-physical-warfare-in-the-middle-east/> Accessed September 28, 2026.
- 105 CISA Adds Five Known Exploited Vulnerabilities to Catalog. Cybersecurity and Infrastructure Security Agency, March 5, 2026.  
<https://www.cisa.gov/news-events/alerts/2026/03/05/cisa-adds-five-known-exploited-vulnerabilities-catalog> Accessed September 28, 2026.
- 106 Investigation into International "ATM Jackpotting" Scheme and Tren de Aragua results in Additional Indictment and 87 Total Charged Defendants. US Attorney's Office, District of Nebraska, January 26, 2026.  
<https://www.justice.gov/usao-ne/pr/investigation-international-atm-jackpotting-scheme-and-tren-de-aragua-results-additional> Accessed September 28, 2026.
- 107 UNC2891 Bank Heist: Physical ATM Backdoor and Linux Forensic Evasion. Group-IB, July 30, 2025. <https://www.group-ib.com/blog/unc2891-bank-heist/> Accessed September 28, 2026.
- 108 Cyber-Attack Against Ukrainian Critical Infrastructure (IR-ALERT-H-16-056-01). CISA (ICS-CERT), February 25, 2016.  
<https://www.cisa.gov/news-events/ics-alerts/ir-alert-h-16-056-01> Accessed September 28, 2026.
- 109 Win32/Industroyer: A new threat for industrial control systems. ESET, June 12, 2017.  
[https://www.welivesecurity.com/wp-content/uploads/2017/06/Win32\\_Industroyer.pdf](https://www.welivesecurity.com/wp-content/uploads/2017/06/Win32_Industroyer.pdf) Accessed September 28, 2026.
- 110 CyOTE Case Study: Triton. Idaho National Laboratory, CyOTE program, 2025-12.  
[https://cyote.inl.gov/content/uploads/24/2025/12/CyOTE-Case-Study\\_TRITON.pdf](https://cyote.inl.gov/content/uploads/24/2025/12/CyOTE-Case-Study_TRITON.pdf) Accessed September 28, 2026.
- 111 Treasury Sanctions Russian Government Research Institution Connected to the Triton Malware. U.S. Department of the Treasury, October 23, 2020.  
<https://home.treasury.gov/news/press-releases/sm1162> Accessed September 28, 2026.
- 112 TRITON Actor TTP Profile, Custom Attack Tools, Detections, and ATT&CK Mapping. FireEye Mandiant (Google Cloud), April 10, 2019.  
<https://cloud.google.com/blog/topics/threat-intelligence/triton-actor-ttp-profile-custom-attack-tools-detections> Accessed September 28, 2026.

- 113 Attackers Deploy New ICS Attack Framework "TRITON" and Cause Operational Disruption to Critical Infrastructure. FireEye Mandiant (Google Cloud), December 14, 2017. <https://cloud.google.com/blog/topics/threat-intelligence/attackers-deploy-new-ics-attack-framework-triton> Accessed September 28, 2026.
- 114 Four Russian Government Employees Charged in Two Historical Hacking Campaigns Targeting Critical Infrastructure Worldwide. U.S. Department of Justice, Office of Public Affairs, March 24, 2022. <https://www.justice.gov/archives/opa/pr/four-russian-government-employees-charged-two-historical-hacking-campaigns-targeting-critical> Accessed September 28, 2026.
- 115 IRGC-Affiliated Cyber Actors Exploit PLCs in Multiple Sectors, Including U.S. Water and Wastewater Systems Facilities (AA23-335A, updated). CISA, FBI, NSA, EPA and INCD, December 18, 2024. <https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-335a> Accessed September 28, 2026.
- 116 Energy Sector Incident Report - 29 December 2025. CERT Polska (NASK), January 30, 2026. [https://cert.pl/uploads/docs/CERT\\_Polska\\_Energy\\_Sector\\_Incident\\_Report\\_2025.pdf](https://cert.pl/uploads/docs/CERT_Polska_Energy_Sector_Incident_Report_2025.pdf) Accessed September 28, 2026.
- 117 Russia behind cyber-attack with Europe-wide impact an hour before Ukraine invasion. UK Government (Foreign, Commonwealth and Development Office and NCSC), May 10, 2022. <https://www.gov.uk/government/news/russia-behind-cyber-attack-with-europe-wide-impact-an-hour-before-ukraine-invasion> Accessed September 28, 2026.
- 118 After Action Report: 2025 Statewide Cyber Incident. State of Nevada Governor's Technology Office (prepared with Info-Tech Research Group), October 28, 2025. [https://www.it.nv.gov/siteassets/itnew.nv.gov/content/governance/GTO\\_Statewide\\_Cyber\\_Event\\_AAR\\_Final.pdf](https://www.it.nv.gov/siteassets/itnew.nv.gov/content/governance/GTO_Statewide_Cyber_Event_AAR_Final.pdf) Accessed September 28, 2026.
- 119 House Homeland, China Select Committee Republicans Demand Answers from CCP-Backed Company Operating at U.S. Ports. US House Committee on Homeland Security, March 7, 2024. <https://homeland.house.gov/2024/03/07/house-homeland-china-select-committee-republicans-demand-answers-from-ccp-backed-company-operating-at-u-s-ports-amid-shocking-joint-investigation-findings/> Accessed September 28, 2026.
- 120 U.S. authorities find no definitive evidence of hidden devices in Chinese solar inverters. pv magazine, January 28, 2026. Press report. <https://www.pv-magazine.com/2026/01/28/u-s-authorities-find-no-definitive-evidence-of-hidden-devices-in-chinese-solar-inverters/> Accessed September 28, 2026.
- 121 AWS Health Dashboard, public event log for ME-CENTRAL-1 (UAE) and ME-SOUTH-1 (Bahrain), Region Availability, March 1 to September 15, 2026. Amazon Web Services, September 15, 2026. <https://health.aws.amazon.com/public/current-events> Accessed September 28, 2026.
- 122 Incidents at Abqaiq and Khurais. Saudi Aramco, September 14, 2019. <https://www.aramco.com/en/news-media/news/2019/incidents-at-abqaiq-and-khurais> Accessed September 28, 2026.
- 123 Security Council Press Statement on Terrorist Attacks in United Arab Emirates (SC/14771). United Nations, Meetings Coverage and Press Releases, January 21, 2022. <https://press.un.org/en/2022/sc14771.doc.htm> Accessed September 28, 2026.
- 124 Security Council Adopts Resolution 2817 (2026) Condemning Iran's 'Egregious Attacks' against Neighbours as Middle East Violence Rapidly Escalates (SC/16315). United Nations, Meetings Coverage and Press Releases, March 11, 2026. <https://press.un.org/en/2026/sc16315.doc.htm> Accessed September 28, 2026.
- 125 Implement Peace Deal Now, Speakers Urge, as Security Council Meets in Wake of Resumed Military Confrontation between United States, Iran (SC/16405). United Nations, Meetings Coverage and Press Releases, July 2, 2026. <https://press.un.org/en/2026/sc16405.doc.htm> Accessed September 28, 2026.
- 126 Operational Risk Management Regulation (C 1/2026). Central Bank of the UAE, September 14, 2026. <https://rulebook.centralbank.ae/en/entiresection/7267> Accessed September 28, 2026.
- 127 CBB Rulebook Volume 1, Module OM Operational Risk Management. Central Bank of Bahrain, 2026-05. <https://cbbn.thomsonreuters.com/entiresection/2300930> Accessed September 28, 2026.
- 128 Cloud Computing Regulation. Qatar Central Bank, April 15, 2024. <https://www.qcb.gov.qa/Documents/InformationSecurity/Cloud%20Computing%20Regulation.pdf> Accessed September 28, 2026.
- 129 Data Handling and Protection Regulation. Qatar Central Bank, January 13, 2025. <https://www.qcb.gov.qa/Documents/InformationSecurity/Data%20Handling%20and%20Protection%20Regulation.pdf> Accessed September 28, 2026.
- 130 Critical Systems Cybersecurity Controls (CSCC-1:2019). National Cybersecurity Authority, Saudi Arabia, 2019. [https://cdn.nca.gov.sa/api/files/public/upload/f15af01c-dc59-4281-95e2-03a770655937\\_Critical-Systems-Cybersecurity-Controls.pdf](https://cdn.nca.gov.sa/api/files/public/upload/f15af01c-dc59-4281-95e2-03a770655937_Critical-Systems-Cybersecurity-Controls.pdf) Accessed September 28, 2026.
- 131 Essential Cybersecurity Controls (ECC-2:2024). National Cybersecurity Authority, Saudi Arabia, 2024. [https://cdn.nca.gov.sa/api/files/public/upload/86e09090-44e4-481f-bc28-355673607654\\_ECC--2024-EN.pdf](https://cdn.nca.gov.sa/api/files/public/upload/86e09090-44e4-481f-bc28-355673607654_ECC--2024-EN.pdf) Accessed September 28, 2026.
- 132 Financial Entities Ethical Red-Teaming, Version 1.0. Saudi Central Bank (SAMA), 2019-05. [https://rulebook.sama.gov.sa/sites/default/files/ar\\_net\\_file\\_store/SAMA\\_AR\\_2898\\_VER1.pdf](https://rulebook.sama.gov.sa/sites/default/files/ar_net_file_store/SAMA_AR_2898_VER1.pdf) Accessed September 28, 2026.
- 133 RansomLook API, posts by year, 2021 to 2026, and About page (data licensed CC BY 4.0). RansomLook open-source project, September 28, 2026. <https://www.ransomlook.io/about> Accessed September 28, 2026.
- 134 State of the UAE Cybersecurity Report 2025. UAE Cyber Security Council and CPX, February 17, 2025. <https://info.cpx.net/hubfs/Cpx-Reports-Whitepapers/state-of-the-uae-cybersecurity-report-2025.pdf> Accessed September 28, 2026.
- 135 12 CFR part 225 subpart N, computer-security incident notification (eCFR current). Office of the Federal Register, eCFR, September 28, 2026. <https://www.ecfr.gov/current/title-12/chapter-II/subchapter-A/part-225/subpart-N> Accessed September 28, 2026.
- 136 12 CFR part 53, computer-security incident notification (eCFR current). Office of the Federal Register, eCFR, September 28, 2026. <https://www.ecfr.gov/current/title-12/chapter-I/part-53> Accessed September 28, 2026.
- 137 CBUAE Rulebook, Operational Risk Management Regulation (effective September 14, 2026). Central Bank of the UAE, September 14, 2026. <https://rulebook.centralbank.ae/en/rulebook/operational-risk-management-regulation> Accessed September 28, 2026.
- 138 Cyber Threats and Incidents. National Cybersecurity Authority, Saudi Arabia, June 17, 2025. <https://nca.gov.sa/en/cyber-operations/> Accessed September 28, 2026.
- 139 UAE Cybersecurity Council announces systematic terrorist cyberattacks targeting vital sectors thwarted. Emirates News Agency (WAM), February 21, 2026. <https://www.wam.ae/en/article/byup8x8-uae-cybersecurity-council-announces-systematic> Accessed September 28, 2026.
- 140 Kuwait MoH.: Essential systems restored, active after recent cyberattack. Kuwait News Agency (KUNA), September 25, 2024. <https://www.kuna.net.kw/ArticleDetails.aspx?id=3183935&language=en> Accessed September 28, 2026.
- 141 Cyber Thematic Review 2024. Dubai Financial Services Authority, 2024. [https://www.dfsa.ae/download\\_file/3762/0](https://www.dfsa.ae/download_file/3762/0) Accessed September 28, 2026.

- 142 Qatar marks one year since QNA hack. Qatar Government Communications Office, May 24, 2018. <https://www.gco.gov.qa/en/media-centre/top-news/qatar-marks-one-year-since-qna-hack/> Accessed September 28, 2026.
- 143 Saudi Aramco facing \$50 million cyber extortion over leaked data (Associated Press). CNBC / Associated Press, July 22, 2021. Press report. <https://www.cnbc.com/2021/07/22/saudi-aramco-facing-50m-cyber-extortion-over-leaked-data.html> Accessed September 28, 2026.
- 144 Qatar National Bank Hackers Threaten Data Leak From Second Bank. SecurityWeek, May 4, 2016. Press report. <https://www.securityweek.com/qatar-national-bank-hackers-could-leak-data-second-bank/> Accessed September 28, 2026.
- 145 Leader Of Three Worldwide Cyberattacks Sentenced To 8 Years For Computer Intrusion And Access Device Fraud Conspiracies. U.S. Attorney's Office, Eastern District of New York, February 10, 2017. <https://www.justice.gov/usao-edny/pr/leader-three-worldwide-cyberattacks-sentenced-8-years-computer-intrusion-and-access> Accessed September 28, 2026.
- 146 Global Navigation Satellite System (GNSS) Outages and Alterations, impacted FIRs in the last 30 days on 23/09/2026. European Union Aviation Safety Agency, September 23, 2026. <https://www.easa.europa.eu/en/domains/air-operations/global-navigation-satellite-system-outages-and-alterations> Accessed September 28, 2026.
- 147 JMIC Advisory Note 001-26, Update 082. Joint Maritime Information Center, via UKMTO, August 9, 2026. <https://www.ukmto.org/-/media/ukmto/products/update-082-jmic-advisory-note-09-august.pdf?rev=3f264c5bd28247a0a6ce7f4fad49cfce> Accessed September 28, 2026.
- 148 FRO: Front Eagle Incident. Frontline plc, June 17, 2025. <https://www.frontlineplc.cy/fro-front-eagle-incident/> Accessed September 28, 2026.
- 149 Security incident disclosure, July 2026. Hugging Face, July 16, 2026. <https://huggingface.co/blog/security-incident-july-2026> Accessed September 28, 2026.
- 150 Anatomy of a Frontier Lab Agent Intrusion: A Technical Timeline of the July 2026 Incident. Hugging Face, July 27, 2026. <https://huggingface.co/blog/agent-intrusion-technical-timeline> Accessed September 28, 2026.
- 151 OpenAI and Hugging Face partner to address security incident during model evaluation. OpenAI, July 21, 2026. <https://openai.com/index/hugging-face-model-evaluation-security-incident/> Accessed September 28, 2026.
- 152 The Hugging Face incident and the road ahead. OpenAI, August 26, 2026. <https://openai.com/index/hugging-face-incident-and-the-road-ahead/> Accessed September 28, 2026.
- 153 Investigating three real-world incidents in our cybersecurity evaluations. Anthropic, July 30, 2026. <https://www.anthropic.com/news/investigating-incidents-cybersecurity-evals> Accessed September 28, 2026.
- 154 Incident Report: unsanctioned agent behaviour during cyber testing. UK AI Security Institute, August 4, 2026. <https://www.aisi.gov.uk/blog/incident-report-unsanctioned-agent-behaviour-during-cyber-testing> Accessed September 28, 2026.
- 155 Security Incident INC-2026-07-28-01, preliminary technical report. UK AI Security Institute, August 4, 2026. [https://cdn.prod.website-files.com/663bd486c5e4c81588db7a1d/6a724858f7db25c81487016d\\_Security%20Incident%20INC-2026-07-28-01.pdf](https://cdn.prod.website-files.com/663bd486c5e4c81588db7a1d/6a724858f7db25c81487016d_Security%20Incident%20INC-2026-07-28-01.pdf) Accessed September 28, 2026.
- 156 Staying ahead of threat actors in the age of AI. Microsoft Threat Intelligence, February 14, 2024. <https://www.microsoft.com/en-us/security/blog/2024/02/14/staying-ahead-of-threat-actors-in-the-age-of-ai/> Accessed September 28, 2026.
- 157 Disrupting malicious uses of AI by state-affiliated threat actors. OpenAI, February 14, 2024. <https://openai.com/index/disrupting-malicious-uses-of-ai-by-state-affiliated-threat-actors/> Accessed September 28, 2026.
- 158 Detecting and countering misuse of AI: August 2025. Anthropic, August 27, 2025. <https://www.anthropic.com/news/detecting-countering-misuse-aug-2025> Accessed September 28, 2026.
- 159 Disrupting the first reported AI-orchestrated cyber espionage campaign. Anthropic, November 13, 2025. <https://www.anthropic.com/news/disrupting-AI-espionage> Accessed September 28, 2026.
- 160 GTIG AI Threat Tracker: Adversaries Leverage AI for Vulnerability Exploitation, Augmented Operations, and Initial Access. Google Threat Intelligence Group, May 12, 2026. <https://cloud.google.com/blog/topics/threat-intelligence/ai-vulnerability-exploitation-initial-access> Accessed September 28, 2026.
- 161 Detecting and countering misuse of AI: September 2026. Anthropic, 2026-09. <https://www.anthropic.com/threat-intelligence-report-september-2026> Accessed September 28, 2026.
- 162 GTIG AI Threat Tracker: From Prompting to Autonomy, the Evolution of Adversarial AI. Google Threat Intelligence Group, September 9, 2026. <https://cloud.google.com/blog/topics/threat-intelligence/from-prompting-to-autonomy-the-evolution-of-adversarial-ai> Accessed September 28, 2026.
- 163 CrowdStrike Uncovers New Agentic AI Campaigns. CrowdStrike, October 21, 2026. <https://www.crowdstrike.com/en-us/blog/crowdstrike-uncovers-new-agentic-ai-campaigns/> Accessed October 26, 2026.
- 164 Disrupting the first reported AI-orchestrated cyber espionage campaign (full report). Anthropic, 2025-11. <https://assets.anthropic.com/m/ec212e6566a0d47/original/Disrupting-the-first-reported-AI-orchestrated-cyber-espionage-campaign.pdf> Accessed September 28, 2026.
- 165 AI in the Breach: How an Adversary Leveraged AI to Target a Water Utility's OT. Dragos, May 6, 2026. <https://www.dragos.com/blog/ai-assisted-ics-attack-water-utility> Accessed September 28, 2026.
- 166 Cato CTRL Threat Research: WormGPT Variants Powered by Grok and Mixtral. Cato Networks, June 17, 2025. <https://www.catonetworks.com/blog/cato-ctrl-wormgpt-variants-powered-by-grok-and-mixtral/> Accessed September 28, 2026.
- 167 The Devil Reviews Xanthorox: A Criminal-Focused Analysis of the Latest Malicious LLM Offering. Trend Micro, November 11, 2025. <https://www.trendmicro.com/vinfo/us/security/news/cybercrime-and-digital-threats/the-devil-reviews-xanthorox-a-criminal-focused-analysis-of-the-latest-malicious-llm-offering> Accessed September 28, 2026.
- 168 An Update on the State of Criminal AI: Crime as a Service, AI as the Multiplier. Trend Micro, January 28, 2026. <https://www.trendmicro.com/vinfo/us/security/news/cybercrime-and-digital-threats/the-state-of-criminal-ai> Accessed September 28, 2026.
- 169 Criminal AI-as-a-Service in 2026: How the Underground Market Is Operationalizing Cybercrime. Rapid7, June 11, 2026. <https://www.rapid7.com/blog/post/tr-criminal-ai-underground-market-operationalizing-cybercrime-2026/> Accessed September 28, 2026.
- 170 Cyberattacks by UAC-0001 on the security and defence sector using LAMEHUG, a program that uses an LLM (CERT-UA#16039). CERT-UA, July 17, 2025. <https://cert.gov.ua/article/6284730> Accessed September 28, 2026.
- 171 Disrupting malicious uses of AI: June 2025. OpenAI, June 5, 2025. <https://cdn.openai.com/threat-intelligence-reports/5f73af09-a3a3-4a55-992e-069237681620/disrupting-malicious-uses-of-ai-june-2025.pdf> Accessed September 28, 2026.
- 172 How Far Behind the Frontier are Leading Open Weight Models on Cyber?. UK AI Security Institute, July 17, 2026. <https://www.aisi.gov.uk/blog/how-far-behind-the-frontier-are-leading-open-weight-models-on-cyber> Accessed September 28, 2026.

- 173 UK AISI / CAISI Preliminary Assessment of Kimi K3's Cyber Capabilities. NIST CAISI and UK AI Security Institute, July 23, 2026. <https://www.nist.gov/news-events/news/2026/07/uk-aisi-caisi-preliminary-assessment-kimi-k3s-cyber-capabilities> Accessed September 28, 2026.
- 174 AI Cyber Challenge marks pivotal inflection point for cyber defense. DARPA, August 8, 2025. <https://www.darpa.mil/news/2025/aixcc-results> Accessed September 28, 2026.
- 175 Google's latest AI security announcements. Google, 2025-07. <https://blog.google/technology/safety-security/cybersecurity-updates-summer-2025/> Accessed September 28, 2026.
- 176 A Mythos moment? Frontier AI and cyber risk (BIS Bulletin No 129). Bank for International Settlements, July 20, 2026. <https://www.bis.org/publications/bulletin-129-mythos-moment-frontier-ai-and-cyber-risk.pdf> Accessed September 28, 2026.
- 177 The zero-days are numbered. Mozilla, April 21, 2026. <https://blog.mozilla.org/en/firefox/privacy-security/ai-security-zero-day-vulnerabilities/> Accessed October 26, 2026.
- 178 Frontier AI and cyber resilience. Financial Conduct Authority, September 2, 2026. <https://www.fca.org.uk/publications/multi-firm-reviews/frontier-ai-cyber-resilience> Accessed October 26, 2026.
- 179 Usage Statistics and Market Share of PHP for Websites, October 2026. W3Techs, October 26, 2026. <https://w3techs.com/technologies/details/pl-php> Accessed October 26, 2026.
- 180 Amazon Threat Intelligence identifies PHP2Shell exploitation prior to public disclosure in MadPot. Amazon Web Services, October 21, 2026. <https://aws.amazon.com/blogs/security/amazon-threat-intelligence-identifies-php2shell-exploitation-prior-to-public-disclosure-in-madpot/> Accessed October 26, 2026.
- 181 Wednesday, October 21, 2026 Security Releases. Node.js, October 21, 2026. <https://nodejs.org/en/blog/vulnerability/october-2026-security-releases> Accessed October 26, 2026.
- 182 Microsoft Security Advisory CVE-2026-111497, .NET Brotli Decoder Out-of-bounds Write Vulnerability. Microsoft, dotnet/runtime, October 25, 2026. <https://github.com/dotnet/runtime/security/advisories/GHSA-p3jp-9v83-76w9> Accessed October 26, 2026.
- 183 M365 Copilot Information Disclosure Vulnerability, CVE-2025-32711. Microsoft Security Response Center, June 11, 2025. <https://msrc.microsoft.com/update-guide/vulnerability/CVE-2025-32711> Accessed September 28, 2026.
- 184 NVD CVE API 2.0 and JSON 2.0 feeds, keyword and full-text counts. National Institute of Standards and Technology, September 28, 2026. <https://services.nvd.nist.gov/rest/json/cves/2.0> Accessed September 28, 2026.
- 185 Malicious versions of Nx were published (GHSA-cxm3-wv7p-598c). GitHub Advisory Database, August 27, 2025. <https://github.com/advisories/GHSA-cxm3-wv7p-598c> Accessed September 28, 2026.
- 186 sIngularity postmortem. Nx, 2025. <https://nx.dev/blog/sIngularity-postmortem> Accessed September 28, 2026.
- 187 Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident. METR with Redwood Research, August 26, 2026. <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/> Accessed September 28, 2026.
- 188 Careful adoption of agentic AI services. ASD's ACSC, CISA, NSA, Canadian Centre for Cyber Security, NCSC-NZ and NCSC-UK, May 1, 2026. [https://media.defense.gov/2026/Apr/30/2003922823/-1/-1/0/CAREFUL%20ADOPTION%20F%20AGENTIC%20AI%20SERVICES\\_FINAL.PDF](https://media.defense.gov/2026/Apr/30/2003922823/-1/-1/0/CAREFUL%20ADOPTION%20F%20AGENTIC%20AI%20SERVICES_FINAL.PDF) Accessed September 28, 2026.
- 189 Guidance Note on the Consumer Protection and Responsible Adoption and Use of Artificial Intelligence and Machine Learning by Licensed Financial Institutions in the U.A.E. Central Bank of the UAE, February 11, 2026. <https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence> Accessed September 28, 2026.
- 190 Artificial Intelligence Guideline. Qatar Central Bank, September 4, 2024. [https://www.qcb.gov.qa/Services/Financial%20Technology/QCB\\_Artificial\\_Intelligence\\_Guideline.pdf](https://www.qcb.gov.qa/Services/Financial%20Technology/QCB_Artificial_Intelligence_Guideline.pdf) Accessed September 28, 2026.
- 191 Draft of AI Cybersecurity Guidelines (AICG-1:2026), public consultation draft. National Cybersecurity Authority, Saudi Arabia, 2026-07. [https://cdn.nca.gov.sa/api/files/public/upload/9ae38ab1-64b4-498e-ab30-f054d3a88adb\\_AI\\_Cybersecurity\\_Controls\\_EN.pdf](https://cdn.nca.gov.sa/api/files/public/upload/9ae38ab1-64b4-498e-ab30-f054d3a88adb_AI_Cybersecurity_Controls_EN.pdf) Accessed September 28, 2026.
- 192 Sound Practices for Responsible Adoption of Artificial Intelligence (AI), consultation report. Financial Stability Board, June 10, 2026. <https://www.fsb.org/uploads/P100626.pdf> Accessed September 28, 2026.
- 193 Supervisory Guidance on Model Risk Management (SR 26-2 attachment). Federal Reserve, FDIC and OCC, April 17, 2026. <https://www.federalreserve.gov/supervisionreg/srletters/SR2602a1.pdf> Accessed September 28, 2026.
- 194 CREST's principles for AI-enabled activities. CREST, May 22, 2026. <https://www.crest-approved.org/crests-principles-for-ai-enabled-activities/> Accessed September 28, 2026.
- 195 CREST announces first cohort of companies to achieve new AI-Enabled Penetration Testing accreditation. CREST, September 3, 2026. <https://www.crest-approved.org/crest-first-ai-pt-accredited-members/> Accessed September 28, 2026.
- 196 CREST adds AI checks to cybersecurity accreditation. IT Europa, July 29, 2026. Press report. <https://www.iteuropa.com/news/crest-adds-ai-checks-cybersecurity-accreditation> Accessed September 28, 2026.
- 197 UAE foils AI-powered 'terrorist cyber attacks' on vital sectors. The National, February 22, 2026. Press report. <https://www.thenationalnews.com/news/uae/2026/02/22/uae-foils-ai-powered-terrorist-cyberattacks-on-vital-sectors/> Accessed September 28, 2026.
- 198 Obstacles to Weaponizing PHP2Shell Likely Constrain Immediate Widespread Exploitation. CrowdStrike, October 23, 2026. <https://www.crowdstrike.com/en-us/blog/obstacles-to-weaponizing-php2shell-likely-constrain-exploitation/> Accessed October 26, 2026.
- 199 Threat Brief: Critical Vulnerabilities in PHP, CVE-2026-109618 and CVE-2026-109619 (updated October 25). Palo Alto Networks Unit 42, October 25, 2026. <https://unit42.paloaltonetworks.com/cve-2026-109618-cve-2026-109619/> Accessed October 26, 2026.
- 200 Withdrawal of Proposed Regulatory Actions, Release No. 33-11377 (90 FR 25531). U.S. Securities and Exchange Commission, June 12, 2025. <https://www.sec.gov/files/rules/final/2025/33-11377.pdf> Accessed September 28, 2026.
- 201 Technology Risk Instructions for Financial Services Operators, V1.2. Qatar Central Bank, June 9, 2022. <https://www.qcb.gov.qa/Documents/Information%20Security/Technology%20Risk%20Instructions%20for%20Financial%20Services%20Operators.pdf> Accessed September 28, 2026.
- 202 Retail Payment Services and Card Schemes Regulation, Article (13) Technology Risk and Information Security. Central Bank of the UAE, June 6, 2021. <https://rulebook.centralbank.ae/en/rulebook/article-13-technology-risk-and-information-security> Accessed September 28, 2026.
- 203 DRAFT FedRAMP Penetration Test Guidance, Version 4.0. FedRAMP Program Management Office, GSA, March 4, 2024. [https://www.fedramp.gov/resources/documents/CSP\\_Penetration\\_Test\\_Guidance\\_public\\_comment.pdf](https://www.fedramp.gov/resources/documents/CSP_Penetration_Test_Guidance_public_comment.pdf) Accessed September 28, 2026.
- 204 23 NYCRR Part 500 Cybersecurity Regulation (consolidated text of the Second Amendment). New York State Department of Financial Services, November 1, 2023. <https://www.dfs.ny.gov/system/files/documents/2026/07/NYCRR-part-500-Cybersecurity-Regulation.pdf> Accessed September 28, 2026.

- 205 16 CFR part 314, Standards for Safeguarding Customer Information, eCFR. Federal Trade Commission, Office of the Federal Register, September 24, 2026.  
<https://www.ecfr.gov/current/title-16/part-314> Accessed September 28, 2026.
- 206 FFIEC Cybersecurity Assessment Tool. Federal Financial Institutions Examination Council, 2017-05.  
[https://www.ffiec.gov/sites/default/files/media/resources/FFIEC\\_CAT\\_May\\_2017.pdf](https://www.ffiec.gov/sites/default/files/media/resources/FFIEC_CAT_May_2017.pdf) Accessed September 28, 2026.
- 207 CAT Sunset Statement. Federal Financial Institutions Examination Council, August 29, 2024.  
<https://www.ffiec.gov/sites/default/files/media/press-releases/2024/cat-sunset-statement-ffiec-letterhead.pdf> Accessed September 28, 2026.
- 208 CFTC Issues Withdrawal of Operational Resilience Framework Proposed Rules, Release 9117-25. Commodity Futures Trading Commission, September 9, 2025.  
<https://www.cftc.gov/PressRoom/PressReleases/9117-25> Accessed September 28, 2026.
- 209 CBEST Implementation Guide. Bank of England and Prudential Regulation Authority, February 23, 2024.  
<https://www.bankofengland.co.uk/-/media/boe/files/financial-stability/financial-sector-continuity/cbest-implementation-guide.pdf> Accessed September 28, 2026.
- 210 TIBER-EU Framework: How to implement the European framework for Threat Intelligence-Based Ethical Red teaming. European Central Bank, 2025-01.  
[https://www.ecb.europa.eu/pub/pdf/other/ecb\\_tiber\\_eu\\_framework\\_2025-b32eff9a10.en.pdf](https://www.ecb.europa.eu/pub/pdf/other/ecb_tiber_eu_framework_2025-b32eff9a10.en.pdf) Accessed September 28, 2026.
- 211 Supervisory Policy Manual TM-C-1, Supervisory Approach on Cyber Risk Management, V.1. Hong Kong Monetary Authority, November 29, 2024. <https://brdr.hkma.gov.hk/eng/doc-ldg/docld/getPdf/20241202-2-EN/TM-C-1.pdf> Accessed September 28, 2026.
- 212 Cybersecurity Fortification Initiative 2.0 (circular). Hong Kong Monetary Authority, November 3, 2020.  
<https://brdr.hkma.gov.hk/eng/doc-ldg/docld/getPdf/20201103-2-EN/20201103-2-EN.pdf> Accessed September 28, 2026.
- 213 Red Team: Adversarial Attack Simulation Exercise Guidelines for the Financial Industry in Singapore, v2.0. The Association of Banks in Singapore, 2024-09.  
<https://www.abs.org.sg/docs/library/abs-red-team-adversarial-attack-simulation-exercises-guidelines----september-2024.pdf> Accessed September 28, 2026.
- 214 CORIE Framework Program Guide for Financial Institutions (including FMI) in Australia, Version 2.0. Council of Financial Regulators, 2022-07.  
<https://www.cfr.gov.au/publications/policy-statements-and-other-reports/2022/revised-corie-framework-rollout/pdf/corie-framework.pdf> Accessed September 28, 2026.
- 215 17 CFR 242.1000 to 242.1007 (Regulation SCI), eCFR. U.S. Securities and Exchange Commission, Office of the Federal Register, September 24, 2026.  
<https://www.ecfr.gov/current/title-17/chapter-II/part-242> Accessed September 28, 2026.
- 216 17 CFR 39.18 (DCO system safeguards), eCFR. Commodity Futures Trading Commission, Office of the Federal Register, September 24, 2026. <https://www.ecfr.gov/current/title-17/section-39.18> Accessed September 28, 2026.
- 217 Cyber Security and Resilience Framework (CS&RF). Central Bank of Oman, 2023-09.  
[https://cbo.gov.om/sites/assets/FintechCompulsoryDocs/Cyber%20Security%20%20Resilience%20Framework%20\(CS&RF\).pdf](https://cbo.gov.om/sites/assets/FintechCompulsoryDocs/Cyber%20Security%20%20Resilience%20Framework%20(CS&RF).pdf) Accessed September 28, 2026.
- 218 Registration and Licensing. National Cybersecurity Authority (Saudi Arabia), August 16, 2026. <https://nca.gov.sa/en/registration-and-licensing/> Accessed September 28, 2026.
- 219 Accredited Service Providers register. National Cyber Security Agency (Qatar), September 28, 2026.  
<https://ncsa.gov.qa/en/services/accreditation-status> Accessed September 28, 2026.
- 220 2025 CBEST thematic. Bank of England, PRA and FCA, January 20, 2026.  
<https://www.bankofengland.co.uk/financial-stability/operational-resilience-of-the-financial-sector/2025-cbest-thematic> Accessed September 28, 2026.
- 221 NCSC Decision No. (36) of 2026. National Cyber Security Center, Bahrain, August 19, 2026.  
[https://www.ncsc.gov.bh/static\\_assets/documents/decision-36-2026.pdf](https://www.ncsc.gov.bh/static_assets/documents/decision-36-2026.pdf) Accessed September 28, 2026.
- 222 Cyber-Security Centre of Excellence: Cyber Wargames. Central Bank of the UAE.  
<https://www.centralbank.ae/en/our-operations/risk-management/cyber-security-centre-of-excellence-1/cyber-wargames/> Accessed September 28, 2026.
- 223 CBK Launches the Cyber and Operational Resilience Framework. Central Bank of Kuwait, December 3, 2025.  
<https://www.cbk.gov.kw/en/cbk-news/announcements-and-press-releases/press-releases/2025/12/202512031000-cbk-launches-the-cyber-operational-resilience-framework-for-local-banks-and> Accessed September 28, 2026.
- 224 FSRA General Rulebook GEN 3.5 Cyber Risk Management. ADGM Financial Services Regulatory Authority, September 16, 2026.  
<https://en.adgm.thomsonreuters.com/entiresection/28467> Accessed September 28, 2026.
- 225 DFSA Rulebook GEN 5.5 Cyber risk management. Dubai Financial Services Authority, June 21, 2023.  
<https://dfsae.thomsonreuters.com/entiresection/26146> Accessed September 28, 2026.
- 226 CBB letter EDS/NA/C/03/2025, Cyber Security Requirements. Central Bank of Bahrain, July 16, 2025.  
[https://cbben.thomsonreuters.com/sites/default/files/net\\_file\\_store/EDS\\_Cyber\\_Security\\_Requirements\\_17\\_July\\_2025.pdf](https://cbben.thomsonreuters.com/sites/default/files/net_file_store/EDS_Cyber_Security_Requirements_17_July_2025.pdf) Accessed September 28, 2026.
- 227 SAMA Cyber Security Framework, 3.3.15 Cyber Security Incident Management. Saudi Central Bank (SAMA), May 24, 2017.  
<https://rulebook.sama.gov.sa/en/3315-cyber-security-incident-management-0> Accessed September 28, 2026.
- 228 SAMA Information Technology Governance Framework, 3.3.8 IT Incident Management. Saudi Central Bank (SAMA), November 4, 2021.  
<https://rulebook.sama.gov.sa/en/338-it-incident-management> Accessed September 28, 2026.
- 229 Implementing Regulation of the Personal Data Protection Law. Saudi Data and AI Authority (SDAIA), 2023-09.  
<https://sdaia.gov.sa/en/SDAIA/about/Documents/ImplementingRegulationPersonalDataProtectionLaw.pdf> Accessed September 28, 2026.
- 230 Order No. (43) of 2022 on technical and organizational measures. Personal Data Protection Authority, Bahrain, March 17, 2022.  
[https://www.pdp.gov.bh/en/assets/pdf/executive-decisions/eng/the\\_be\\_met\\_in\\_the\\_technical.pdf](https://www.pdp.gov.bh/en/assets/pdf/executive-decisions/eng/the_be_met_in_the_technical.pdf) Accessed September 28, 2026.
- 231 CBUAE Consumer Protection Standards, Article 6: Protection of Consumer Data and Assets. Central Bank of the UAE, 2021.  
<https://rulebook.centralbank.ae/en/rulebook/article-6-protection-consumer-data-and-assets> Accessed September 28, 2026.
- 232 CBB Rulebook Volume 1, OM-5.5.58 (Section B report and weekly updates). Central Bank of Bahrain, 2022-04.  
<https://cbben.thomsonreuters.com/rulebook/om-5558> Accessed September 28, 2026.
- 233 Advisory on Ransomware and the Use of the Financial System to Facilitate Ransom Payments (FIN-2021-A004). Financial Crimes Enforcement Network, November 8, 2021.  
[https://www.fincen.gov/sites/default/files/advisory/2021-11-08/FinCEN%20Ransomware%20Advisory\\_FIN\\_AL\\_508\\_.pdf](https://www.fincen.gov/sites/default/files/advisory/2021-11-08/FinCEN%20Ransomware%20Advisory_FIN_AL_508_.pdf) Accessed September 28, 2026.
- 234 31 CFR 1020.320, Reports by banks of suspicious transactions (eCFR current). Office of the Federal Register, eCFR, September 28, 2026. <https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1020/subpart-C/section-1020.320> Accessed September 28, 2026.
- 235 17 CFR 248.30, Regulation S-P procedures to safeguard customer information (eCFR current). Office of the Federal Register, eCFR, September 28, 2026.  
<https://www.ecfr.gov/current/title-17/chapter-II/part-248/subpart-A/section-248.30> Accessed September 28, 2026.

- 236 Enhancements to Regulation S-P: A Small Entity Compliance Guide. US Securities and Exchange Commission, 2024. <https://www.sec.gov/files/rules/final/2024/regulation-s-p-small-entity-compliance-guide.pdf> Accessed September 28, 2026.
- 237 Withdrawal of Proposed Regulatory Actions (FR Doc. 2025-11110). US Securities and Exchange Commission, Federal Register, June 17, 2025. <https://www.govinfo.gov/content/pkg/FR-2025-06-17/pdf/2025-11110.pdf> Accessed September 28, 2026.
- 238 ADGM's FSRA Issues Cyber Risk Management Framework. ADGM Financial Services Regulatory Authority, July 29, 2025. <https://www.adgm.com/media/announcements/adgms-fsra-issues-cyber-risk-management-framework> Accessed September 28, 2026.
- 239 California Civil Code 1798.82, as amended by Stats. 2025, Ch. 319 (SB 446). California Legislative Information, 2025. [https://leginfo.ca.gov/faces/codes\\_displaySection.xhtml?lawCode=CIV&sectionNum=1798.82](https://leginfo.ca.gov/faces/codes_displaySection.xhtml?lawCode=CIV&sectionNum=1798.82) Accessed September 28, 2026.
- 240 6 U.S.C. 681b, Required reporting of certain cyber incidents. US Government Publishing Office, 2023. <https://www.govinfo.gov/content/pkg/USCODE-2023-title6/html/USCODE-2023-title6-chap1-subchapXVIII-partD-sec681b.htm> Accessed September 28, 2026.
- 241 Unified Agenda entry RIN 1670-AA04, CIRCIA Reporting Requirements. Office of Information and Regulatory Affairs, 2026. <https://www.reginfo.gov/public/do/eAgendaViewRule?pubId=202510&RIN=1670-AA04> Accessed September 28, 2026.
- 242 Federal Register API search for CIRCIA documents. Office of the Federal Register, September 28, 2026. [https://www.federalregister.gov/api/v1/documents.json?conditions\[term\]=CIRCIA&order=newest](https://www.federalregister.gov/api/v1/documents.json?conditions[term]=CIRCIA&order=newest) Accessed September 28, 2026.
- 243 Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA). Cybersecurity and Infrastructure Security Agency, September 28, 2026. <https://www.cisa.gov/topics/cyber-threats-and-advisories/information-sharing/cyber-incident-reporting-critical-infrastructure-act-2022-circia> Accessed September 28, 2026.
- 244 Executive Regulation of the Personal Data Protection Law, Ministerial Decision 34/2024. Ministry of Transport, Communications and Information Technology, Oman, January 28, 2024. <https://mtcit.gov.om/web/content/41508?unique=ea5b8bcd9c6d2ae1be960873fddd272ce22a778&download=true> Accessed September 28, 2026.
- 245 ADGM Data Protection Regulations 2021, section 32, Notification of a Personal Data Breach to the Commissioner of Data Protection. Abu Dhabi Global Market, 2021. <https://en.adgm.thomsonreuters.com/rulebook/32-notification-personal-data-breach-commissioner-data-protection> Accessed September 28, 2026.
- 246 QFC Data Protection Regulations 2021, Article 31, Notification of personal data breaches. Qatar Financial Centre, December 6, 2023. <https://qfca-en.thomsonreuters.com/rulebook/article-31-notification-personal-data-breaches> Accessed September 28, 2026.
- 247 Regulation (EU) 2016/679 (General Data Protection Regulation). Official Journal of the European Union, May 4, 2016. <https://eur-lex.europa.eu/eli/reg/2016/679/oj> Accessed September 28, 2026.
- 248 Personal Data Breach Notifications: Guidelines for Regulated Entities (PDPPL-02050217E v2.0). National Cyber Security Agency, Qatar, 2022-09. [https://ncsa.gov.qa/shared/document\\_library/94787172324409/33000/8870/1.0~740b1248-e3a9-40e1-b66e-589091cc182a?mimetype=application%2Fpdf&filename=Personal+Data+Breach+Notifications+-+Guideline+for+Regulated+Entities+%284%29.pdf](https://ncsa.gov.qa/shared/document_library/94787172324409/33000/8870/1.0~740b1248-e3a9-40e1-b66e-589091cc182a?mimetype=application%2Fpdf&filename=Personal+Data+Breach+Notifications+-+Guideline+for+Regulated+Entities+%284%29.pdf) Accessed September 28, 2026.
- 249 Federal Decree-Law No. 45 of 2021 on the Protection of Personal Data. United Arab Emirates, September 20, 2021. <https://uaelegislation.gov.ae/en/legislations/1972> Accessed September 28, 2026.
- 250 RFC 7489: Domain-based Message Authentication, Reporting, and Conformance (DMARC). IETF, March 1, 2015. <https://www.rfc-editor.org/rfc/rfc7489.txt> Accessed September 28, 2026.
- 251 RFC 7208: Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1. IETF, April 1, 2014. <https://www.rfc-editor.org/rfc/rfc7208.txt> Accessed September 28, 2026.
- 252 RFC 9116: A File Format to Aid in Security Vulnerability Disclosure. IETF, April 1, 2022. <https://www.rfc-editor.org/rfc/rfc9116.txt> Accessed September 28, 2026.
- 253 BOD 20-01: Develop and Publish a Vulnerability Disclosure Policy. Cybersecurity and Infrastructure Security Agency, September 2, 2020. <https://www.cisa.gov/news-events/directives/bod-20-01-develop-and-publish-vulnerability-disclosure-policy> Accessed September 28, 2026.
- 254 dotgov-data current-full.csv (commit 60930e22948de65a6caa324f8e24d181b00df8b). Cybersecurity and Infrastructure Security Agency, September 28, 2026. <https://raw.githubusercontent.com/cisagov/dotgov-data/main/current-full.csv> Accessed September 28, 2026.
- 255 RFC 8461: SMTP MTA Strict Transport Security (MTA-STS). IETF, September 1, 2018. <https://www.rfc-editor.org/rfc/rfc8461.txt> Accessed September 28, 2026.
- 256 Large Holding Companies (assets as of June 30, 2026). Federal Financial Institutions Examination Council, National Information Center, June 30, 2026. <https://www.ffiec.gov/npw/Institution/TopHoldings> Accessed September 28, 2026.
- 257 Readout: Financial Stability Oversight Council Meeting on July 15, 2026. U.S. Department of the Treasury, July 15, 2026. <https://home.treasury.gov/news/press-releases/sb0565> Accessed September 28, 2026.
- 258 Licensed banks. Saudi Central Bank, September 28, 2026. [https://www.sama.gov.sa/en-us/\\_LAYOUTS/15/SAMA.Portal/PortalHandler.ashx?op=LoadItems&listUrl=/ar-sa/Supervision/LicenseEntities/Lists/LicensedBanks&viewName=Archive](https://www.sama.gov.sa/en-us/_LAYOUTS/15/SAMA.Portal/PortalHandler.ashx?op=LoadItems&listUrl=/ar-sa/Supervision/LicenseEntities/Lists/LicensedBanks&viewName=Archive) Accessed September 28, 2026.
- 259 CB Register as of August 2026. Central Bank of the UAE. <https://centralbank.ae/media/hw0lreyi/cb-register-august-2026.pdf> Accessed September 28, 2026.
- 260 Useful links: Local Banks. Qatar Central Bank, September 28, 2026. [https://www.qcb.gov.qa/\\_api/web/lists/getbytitle\('UsefulLinks'\)/items?\\$select=ID,Title,TitleAR,Link,LinkAR,Modified,Category/Title&\\$expand=Category&\\$top=300](https://www.qcb.gov.qa/_api/web/lists/getbytitle('UsefulLinks')/items?$select=ID,Title,TitleAR,Link,LinkAR,Modified,Category/Title&$expand=Category&$top=300) Accessed September 28, 2026.
- 261 Regulated entities: Kuwaiti banks. Central Bank of Kuwait, September 28, 2026. <https://www.cbk.gov.kw/en/supervision/regulate-d-entities/kuwaiti-banks/conventional-banks> Accessed September 28, 2026.
- 262 CBB Register (licensing directory). Central Bank of Bahrain, September 28, 2026. <https://www.cbb.gov.bh/licensing-directory/> Accessed September 28, 2026.
- 263 Licensed Banks. Central Bank of Oman, September 28, 2026. <https://cbo.gov.om/Pages/LicensedBanks.aspx> Accessed September 28, 2026.
- 264 North Korean Regime-Backed Programmer Charged With Conspiracy to Conduct Multiple Cyber Attacks and Intrusions. U.S. Department of Justice, Office of Public Affairs, September 6, 2018. <https://www.justice.gov/archives/opa/pr/north-korean-regime-backed-programmer-charged-conspiracy-conduct-multiple-cyber-attacks-and> Accessed September 28, 2026.
- 265 Hosted Exchange issues (incident status updates, final update 2023-01-05). Rackspace Technology, January 5, 2023. <https://status.apps.rackspace.com/index/viewincidents?group=2> Accessed September 28, 2026.

- 266 Microsoft Corporation Form 8-K (Item 1.05). Microsoft / US SEC EDGAR, January 19, 2024. <https://www.sec.gov/Archives/edgar/data/789019/000119312524011295/d708866d8k.htm> Accessed September 28, 2026.
- 267 North Korean Regime-Backed Programmer Charged With Conspiracy to Conduct Multiple Cyber Attacks and Intrusions. U.S. Department of Justice, Office of Public Affairs, September 6, 2018. <https://www.justice.gov/archives/opa/pr/north-korean-regime-backed-programmer-charged-conspiracy-conduct-multiple-cyber-attacks-and> Accessed September 28, 2026.
- 268 Unauthorized Access to Okta's Support Case Management System: Root Cause and Remediation. Okta Security, November 3, 2023. <https://sec.okta.com/articles/2023/11/unauthorized-access-oktas-support-case-management-system-root-cause> Accessed September 28, 2026.
- 269 Leader Of Three Worldwide Cyberattacks Sentenced To 8 Years For Computer Intrusion And Access Device Fraud Conspiracies. U.S. Attorney's Office, Eastern District of New York, February 10, 2017. <https://www.justice.gov/usao-edny/pr/leader-three-worldwide-cyberattacks-sentenced-8-years-computer-intrusion-and-access> Accessed September 28, 2026.
- 270 Cyber mafias strike at leading Omani bank. Gulf News, 2013-02. Press report. <https://gulfnews.com/news/gulf/oman/cyber-mafias-strike-at-leading-omani-bank-1.1151049> Accessed September 28, 2026.
- 271 The Great Bank Robbery: the Carbanak APT. Kaspersky Lab, Securelist, February 16, 2015. <https://securelist.com/the-great-bank-robbery-the-carbanak-apt/68732/> Accessed September 28, 2026.
- 272 Member of Sophisticated China-Based Hacking Group Indicted for Series of Computer Intrusions, Including 2015 Data Breach of Health Insurer Anthem Inc. Affecting Over 78 Million People. U.S. Department of Justice, Office of Public Affairs, May 9, 2019. <https://www.justice.gov/opa/pr/member-sophisticated-china-based-hacking-group-indicted-series-computer-intrusions-including> Accessed September 28, 2026.
- 273 United States v. Park Jin Hyok, criminal complaint and affidavit (C.D. Cal., filed 2018-06-08). U.S. Department of Justice, June 8, 2018. <https://www.justice.gov/opa/press-release/file/1092091/download> Accessed September 28, 2026.
- 274 The Home Depot, Inc. Form 10-Q for the quarter ended November 2, 2014. U.S. Securities and Exchange Commission (EDGAR), 2014-11. [https://www.sec.gov/Archives/edgar/data/0000354950/000035495014000047/hd\\_10qx11022014.htm](https://www.sec.gov/Archives/edgar/data/0000354950/000035495014000047/hd_10qx11022014.htm) Accessed September 28, 2026.
- 275 The OPM Data Breach: How the Government Jeopardized Our National Security for More than a Generation, Majority Staff Report. U.S. House Committee on Oversight and Government Reform, September 7, 2016. <https://oversight.house.gov/wp-content/uploads/2016/09/The-OPM-Data-Breach-How-the-Government-Jeopardized-Our-National-Security-for-More-than-a-Generation.pdf> Accessed September 28, 2026.
- 276 United States v. Park Jin Hyok, criminal complaint and affidavit (C.D. Cal., filed 2018-06-08). U.S. Department of Justice, June 8, 2018. <https://www.justice.gov/opa/press-release/file/1092091/download> Accessed September 28, 2026.
- 277 Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace. U.S. Department of Justice, Office of Public Affairs, October 19, 2020. <https://www.justice.gov/archives/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and> Accessed September 28, 2026.
- 278 Qatar National Bank Data Breach. Have I Been Pwned, 2016-04. Press report. <https://haveibeenpwned.com/Breach/QatarNationalBank> Accessed September 28, 2026.
- 279 Qatar National Bank leak: Security experts hint 'SQL injection' used in database hack. IBTimes UK, April 27, 2016. Press report. <https://www.ibtimes.co.uk/qatar-national-bank-leak-security-experts-hint-sql-injection-used-database-hack-1557069> Accessed September 28, 2026.
- 280 The full Shamoon: How the devastating malware was inserted into networks. IBM X-Force IRIS, 2017-01. <https://www.ibm.com/think/x-force/the-full-shamoon-how-the-devastating-malware-was-inserted-into-networks> Accessed September 28, 2026.
- 281 Shamoon 2: Return of the Distrack Wiper. Palo Alto Networks Unit 42, November 30, 2016. <https://unit42.paloaltonetworks.com/unit42-shamoon-2-return-distrack-wiper/> Accessed September 28, 2026.
- 282 Second Wave of Shamoon 2 Attacks Identified. Palo Alto Networks Unit 42, January 9, 2017. <https://unit42.paloaltonetworks.com/unit42-second-wave-shamoon-2-attacks-identified/> Accessed September 28, 2026.
- 283 Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace. U.S. Department of Justice, Office of Public Affairs, October 19, 2020. <https://www.justice.gov/archives/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and> Accessed September 28, 2026.
- 284 The Equifax Data Breach, Majority Staff Report. U.S. House Committee on Oversight and Government Reform, 2018-12. <https://oversight.house.gov/wp-content/uploads/2018/12/Equifax-Report.pdf> Accessed September 28, 2026.
- 285 Chinese Military Personnel Charged with Computer Fraud, Economic Espionage and Wire Fraud for Hacking into Credit Reporting Agency Equifax. U.S. Department of Justice, Office of Public Affairs, February 10, 2020. <https://www.justice.gov/archives/opa/pr/chinese-military-personnel-charged-computer-fraud-economic-espionage-and-wire-fraud-hacking> Accessed September 28, 2026.
- 286 New Ransomware Variant "Nyetya" Compromises Systems Worldwide. Cisco Talos, June 27, 2017. <https://blog.talosintelligence.com/worldwide-ransomware-variant/> Accessed September 28, 2026.
- 287 Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace. U.S. Department of Justice, Office of Public Affairs, October 19, 2020. <https://www.justice.gov/archives/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and> Accessed September 28, 2026.
- 288 How QNA website was hacked? Mol explains. The Peninsula (reporting Qatar Ministry of Interior statements), July 20, 2017. <https://thepeninsulaqatar.com/article/20/07/2017/How-QNA-website-was-hacked-Mol-explains> Accessed September 28, 2026.
- 289 DNSspionage Campaign Targets Middle East. Cisco Talos, November 27, 2018. <https://blog.talosintelligence.com/dnsponage-campaign-targets-middle-east/> Accessed September 28, 2026.
- 290 Global DNS Hijacking Campaign: DNS Record Manipulation at Scale. FireEye Mandiant (Google Cloud), January 10, 2019. <https://cloud.google.com/blog/topics/threat-intelligence/global-dns-hijacking-campaign-dns-record-manipulation-at-scale> Accessed September 28, 2026.
- 291 Emergency Directive 19-01: Mitigate DNS Infrastructure Tampering. Cybersecurity and Infrastructure Security Agency, January 22, 2019. <https://www.cisa.gov/news-events/directives/ed-19-01-mitigate-dns-infrastructure-tampering> Accessed September 28, 2026.
- 292 A Deep Dive on the Recent Widespread DNS Hijacking Attacks. KrebsOnSecurity, February 18, 2019. Press report. <https://krebsonsecurity.com/2019/02/a-deep-dive-on-the-recent-widespread-dns-hijacking-attacks/> Accessed September 28, 2026.

- 293 Dustman wiper attack on Bapco oil company. Kaspersky ICS CERT, January 10, 2020. <https://ics-cert.kaspersky.com/publications/bapco-dustman/> Accessed September 28, 2026.
- 294 Saudi cyber authority uncovers new data-wiping malware, and experts suspect Iran is behind it. CyberScoop, January 8, 2020. Press report. <https://cyberscoop.com/saudi-arabia-iran-cyberattack-soleimani/> Accessed September 28, 2026.
- 295 United States v. Paige Thompson (case summary). US Attorney's Office, Western District of Washington. <https://www.justice.gov/usao-wdwa/united-states-v-paige-thompson> Accessed September 28, 2026.
- 296 Former hacker sentenced for stealing computer power to mine cryptocurrency and stealing the personal information of more than 100 million people. US Attorney's Office, Western District of Washington, October 4, 2022. <https://www.justice.gov/usao-wdwa/pr/former-hacker-sentenced-stealing-computer-power-mine-cryptocurrency-and-stealing> Accessed September 28, 2026.
- 297 Consent Order for the Assessment of a Civil Money Penalty, Capital One, N.A. (AA-EC-20-51, EA 2020-036). Office of the Comptroller of the Currency, 2020-08. <https://www OCC.gov/static/enforcement-actions/ea2020-036.pdf> Accessed September 28, 2026.
- 298 Hackers hit Norsk Hydro with ransomware. The company responded with transparency. Microsoft, December 16, 2019. <https://news.microsoft.com/source/features/digital-transformation/hackers-hit-norsk-hydro-ransomware-company-responded-transparency/> Accessed September 28, 2026.
- 299 Tortoiseshell Group Targets IT Providers in Saudi Arabia in Probable Supply Chain Attacks. Symantec Threat Hunter Team (Broadcom), September 18, 2019. <https://www.security.com/threat-intelligence/tortoiseshell-apt-supply-chain> Accessed September 28, 2026.
- 300 xHunt Campaign: Newly Discovered Backdoors Using Deleted Email Drafts and DNS Tunneling for Command and Control. Palo Alto Networks Unit 42, November 9, 2020. <https://unit42.paloaltonetworks.com/xhunt-campaign-backdoors/> Accessed September 28, 2026.
- 301 xHunt Campaign: New BumbleBee Webshell and SSH Tunnels Used for Lateral Movement. Palo Alto Networks Unit 42, January 11, 2021. <https://unit42.paloaltonetworks.com/bumblebee-webshell-xhunt-campaign/> Accessed September 28, 2026.
- 302 REvil ransomware actors attack Kaseya in supply chain attack. Cisco Talos, July 3, 2021. <https://blog.talosintelligence.com/revil-ransomware-actors-attack-kaseya/> Accessed September 28, 2026.
- 303 HAFNIUM targeting Exchange Servers with 0-day exploits. Microsoft Threat Intelligence, March 2, 2021. <https://www.microsoft.com/en-us/security/blog/2021/03/02/hafnium-targeting-exchange-servers/> Accessed September 28, 2026.
- 304 Cisco Talos shares insights related to recent cyber attack on Cisco. Cisco Talos, August 10, 2022. <https://blog.talosintelligence.com/recent-cyber-attack/> Accessed September 28, 2026.
- 305 Notice of Recent Security Incident. LastPass, December 22, 2022. <https://blog.lastpass.com/posts/notice-of-recent-security-incident> Accessed September 28, 2026.
- 306 Security Incident Update and Recommended Actions. LastPass, March 1, 2023. <https://blog.lastpass.com/posts/security-incident-update-recommended-actions> Accessed September 28, 2026.
- 307 Australian Information Commissioner v Medibank Private Limited: Concise Statement (VID497/2024). Office of the Australian Information Commissioner, filed in the Federal Court of Australia, June 19, 2024. [https://www.oaic.gov.au/\\_data/assets/pdf\\_file/0025/221974/Australian-Information-Commissioner-v-Medibank-Private-Limited-concise-statement.pdf](https://www.oaic.gov.au/_data/assets/pdf_file/0025/221974/Australian-Information-Commissioner-v-Medibank-Private-Limited-concise-statement.pdf) Accessed September 28, 2026.
- 308 Review of the Attacks Associated with Lapsus\$ and Related Threat Groups. Cyber Safety Review Board, US Department of Homeland Security, July 24, 2023. [https://www.cisa.gov/sites/default/files/2023-08/CSRB\\_Lapsus%24\\_508c.pdf](https://www.cisa.gov/sites/default/files/2023-08/CSRB_Lapsus%24_508c.pdf) Accessed September 28, 2026.
- 309 OWASSRF: CrowdStrike Identifies New Exploit Method for Exchange Bypassing ProxyNotShell Mitigations. CrowdStrike, December 20, 2022. <https://www.crowdstrike.com/en-us/blog/owassrf-exploit-analysis-and-recommendations/> Accessed September 28, 2026.
- 310 Rackspace Technology Form 10-K for fiscal year 2023. US Securities and Exchange Commission (EDGAR). <https://www.sec.gov/Archives/edgar/data/181019/00018101924000042/rxt-20231231.htm> Accessed September 28, 2026.
- 311 Security update. Uber, September 16, 2022. <https://www.uber.com/newsroom/security-update/> Accessed September 28, 2026.
- 312 KA-SAT Network cyber attack overview. Viasat. <https://news.viasat.com/blog/corporate/ka-sat-network-cyber-attack-overview> Accessed September 28, 2026.
- 313 Kuwait isolates some government systems following attack on its Finance Ministry. The Record from Recorded Future News, September 26, 2023. Press report. <https://therecord.media/kuwait-isolates-systems-after-ransomware-attack> Accessed September 28, 2026.
- 314 Okta October 2023 Security Incident Investigation Closure. Okta Security, February 8, 2024. <https://sec.okta.com/harfiles> Accessed September 28, 2026.
- 315 Earth Simnavaz (aka APT34) Levies Advanced Cyberattacks Against Middle East. Trend Micro, October 11, 2024. [https://www.trendmicro.com/en\\_us/research/24/j/earth-simnavaz-cyberattacks.html](https://www.trendmicro.com/en_us/research/24/j/earth-simnavaz-cyberattacks.html) Accessed September 28, 2026.
- 316 Iran's APT34 Abuses MS Exchange to Spy on Gulf Gov'ts. Dark Reading, October 17, 2024. Press report. <https://www.darkreading.com/cyberattacks-data-breaches/iran-apt34-ms-exchange-spy-gulf-govts> Accessed September 28, 2026.
- 317 Kuwait Health Ministry restoring systems after cyberattack takes down hospitals, healthcare app. The Record from Recorded Future News, September 26, 2024. Press report. <https://therecord.media/kuwait-ministry-restoring-systems-cyberattack> Accessed September 28, 2026.
- 318 BeyondTrust Remote Support SaaS Service Security Investigation. BeyondTrust, December 8, 2024. <https://www.beyondtrust.com/remote-support-saas-service-security-investigation> Accessed September 28, 2026.
- 319 Letter to Secretary Yellen on CCP cyber breach. Senate Banking Committee Ranking Member and House Financial Services Vice Chair, December 31, 2024. <https://www.banking.senate.gov/imo/media/doc/lettertosecretaryyellenonccp cyberbreach.pdf> Accessed September 28, 2026.
- 320 Joint Statement from FBI and CISA on the People's Republic of China (PRC) Targeting of Commercial Telecommunications Infrastructure. FBI and CISA, November 13, 2024. <https://www.cisa.gov/news-events/news/joint-statement-fbi-and-cisa-peoples-republic-china-prc-targeting-commercial-telecommunications> Accessed September 28, 2026.
- 321 Coinbase breach tied to bribed TaskUs support agents in India. BleepingComputer, June 3, 2025. Press report. <https://www.bleepingcomputer.com/news/security/coinbase-breach-tied-to-bribed-taskus-support-agents-in-india/> Accessed September 28, 2026.
- 322 Collins Aerospace: Old Passwords and Delayed Response Enable Data Theft. heise online, October 26, 2025. Press report. <https://www.heise.de/en/news/Collins-Aerospace-Old-Passwords-and-Delayed-Response-Enable-Data-Theft-10900183.html> Accessed September 28, 2026.
- 323 F5, Inc. Form 8-K, Item 1.05 Material Cybersecurity Incidents. U.S. Securities and Exchange Commission (F5 filing), October 15, 2025. <https://www.sec.gov/Archives/edgar/data/1048695/000104869525000149/f5iv-20251015.htm> Accessed September 28, 2026.

324 ED 26-01: Mitigate Vulnerabilities in F5 Devices. CISA, October 15, 2025. <https://www.cisa.gov/news-events/directives/ed-26-01-mitigate-vulnerabilities-f5-devices> Accessed September 28, 2026.

325 Marks & Spencer breach linked to Scattered Spider ransomware attack. BleepingComputer, April 28, 2025. Press report. <https://www.bleepingcomputer.com/news/security/marks-and-spencer-breach-linked-to-scattered-spider-ransomware-attack/> Accessed September 28, 2026.

326 16 Minutes to Impact: npm Supply Chain Abuse Deploys crypto-draining malware. Sygnia, September 11, 2025. <https://www.sygnia.co/threat-reports-and-advisories/npm-supply-chain-attack-september-2025/> Accessed September 28, 2026.

327 Widespread Supply Chain Compromise Impacting npm Ecosystem. CISA, September 23, 2025. <https://www.cisa.gov/news-events/alerts/2025/09/23/widespread-supply-chain-compromise-impacting-npm-ecosystem> Accessed September 28, 2026.

328 OCC Notifies Congress of Incident Involving Email System (News Release 2025-30). Office of the Comptroller of the Currency, April 8, 2025. <https://www.occ.gov/news-issuances/news-releases/2025/nr-occ-2025-30.html> Accessed September 28, 2026.

329 Disrupting active exploitation of on-premises SharePoint vulnerabilities. Microsoft Threat Intelligence, July 22, 2025. <https://www.microsoft.com/en-us/security/blog/2025/07/22/disrupting-active-exploitation-of-on-premises-sharepoint-vulnerabilities/> Accessed September 28, 2026.

330 UPDATE: Microsoft Releases Guidance on Exploitation of SharePoint Vulnerabilities. CISA, July 20, 2025. <https://www.cisa.gov/news-events/alerts/2025/07/20/update-microsoft-releases-guidance-exploitation-sharepoint-vulnerabilities> Accessed September 28, 2026.

331 Oracle E-Business Suite Zero-Day Exploited in Widespread Extortion Campaign. Google Threat Intelligence Group, October 9, 2025. <https://cloud.google.com/blog/topics/threat-intelligence/oracle-ebusiness-suite-zero-day-exploitation> Accessed September 28, 2026.

332 Drift Security Update (Mandiant investigation update as of 6 September 2025). Salesloft, September 6, 2025. <https://trust.salesloft.com/?uid=Drift+Security+Update> Accessed September 28, 2026.

333 The impact of the Salesloft Drift breach on Cloudflare and our customers. Cloudflare, September 2, 2025. <https://blog.cloudflare.com/response-to-salesloft-drift-incident/> Accessed September 28, 2026.

334 Bitget Says Suspected North Korean Hackers Stole \$351.6M After Backend Compromise. The Hacker News, September 26, 2026. Press report. <https://thehackernews.com/2026/09/bitget-says-suspected-north-korean.html> Accessed September 28, 2026.

335 Bitget resumes Bitcoin withdrawals after \$387.5 million crypto heist. BleepingComputer, September 28, 2026. Press report. <https://www.bleepingcomputer.com/news/security/bitget-resumes-bitcoin-withdrawals-after-3875-million-crypto-heist/> Accessed September 28, 2026.

336 EVERTEC, Inc. Form 8-K (Item 8.01). Evertec / SEC EDGAR, June 9, 2026. <https://www.sec.gov/Archives/edgar/data/0001559865/000155986526000039/evtc-20260609.htm> Accessed September 28, 2026.

337 Energy Sector Incident Report - 29 December 2025. CERT Polska (NASK), January 30, 2026. [https://cert.pl/uploads/docs/CERT\\_Polska\\_Energy\\_Sector\\_Incident\\_Report\\_2025.pdf](https://cert.pl/uploads/docs/CERT_Polska_Energy_Sector_Incident_Report_2025.pdf) Accessed September 28, 2026.

338 Poland Energy Sector Cyber Incident Highlights OT and ICS Security Gaps. CISA, February 10, 2026. <https://www.cisa.gov/news-events/alerts/2026/02/10/poland-energy-sector-cyber-incident-highlights-ot-and-ics-security-gaps> Accessed September 28, 2026.

339 Poland Energy Sector Cyber Incident Highlights OT and ICS Security Gaps. CISA, February 10, 2026. <https://www.cisa.gov/news-events/alerts/2026/02/10/poland-energy-sector-cyber-incident-highlights-ot-and-ics-security-gaps> Accessed September 28, 2026.

340 A hacker accessed sensitive UAE-linked systems then demanded \$5m. Gulf News, September 15, 2026. Press report. <https://gulfnews.com/living-in-uae/safety-security/a-hacker-accessed-sensitive-uae-linked-systems-then-demanded-5m-1.500675569> Accessed September 28, 2026.


341 Stryker Corporation Form 8-K (Item 8.01). Stryker / SEC EDGAR, March 11, 2026. <https://www.sec.gov/Archives/edgar/data/310764/000119312526102460/d76279d8k.htm> Accessed September 28, 2026.

342 Customer Updates: Stryker Network Disruption. Stryker, April 1, 2026. <https://www.stryker.com/us/en/about/news/2026/a-message-to-our-customers-03-2026.html> Accessed September 28, 2026.

343 Justice Department Disrupts Iranian Cyber Enabled Psychological Operations. US Department of Justice, March 19, 2026. <https://www.justice.gov/opa/pr/justice-department-disrupts-iranian-cyber-enabled-psychological-operations> Accessed September 28, 2026.

344 Lessons from the Stryker Incident. Sygnia, March 22, 2026. <https://www.sygnia.co/threat-reports-and-advisories/identity-control-plane-attack-stryker/> Accessed September 28, 2026.

345 FBI seizes Handala data leak site after Stryker cyberattack. BleepingComputer, March 19, 2026. Press report. <https://www.bleepingcomputer.com/news/security/fbi-seizes-handala-data-leak-site-after-stryker-cyberattack/> Accessed September 28, 2026.



# The Seam Report 2026

Bendings is a cybersecurity firm founded in 2013. It runs reconnaissance, red team, penetration testing, physical and social-engineering testing, blue team support and incident response for governments, banks and asset managers in the United States and the Gulf Cooperation Council states. Its own AI model, Bendings Gauge, runs continuous red and blue team work inside limits each client approves, on Bendings Frame, the firm's own operating system for digital and physical security operations.