COMPANION PAPER 1

Evidence and References

The citation record behind the IQ4hire 4D Cybersecurity Framework



IQ4hire Research · November 2026
Proposed and not yet validated

COMPANION PAPER 1

Evidence and References

IQ4hire Research · November 2026 · Proposed and not yet validated

Abstract

The framework paper makes factual claims about attacker speed, patching, third parties and quantum timelines. This paper is the citation record for them. For each source it gives the figure, what the figure measures, what the source does not say, and how far we read it.

1. Method

1. **Primary sources first.** We cite the organization that measured or wrote the thing. A news story is used only to find the primary source and is never the citation for a figure.
2. **Source type is stated.** Most security research comes from firms that sell security products. Each source is marked as vendor, government or standards body, and a vendor figure is cited as that vendor's figure.
3. **Read in full or marked.** Each entry says what we read and what we did not open.
4. **Quote where a number carries the argument.**
5. **State what the figure measures.** A burst against prepared targets is not a disclosure-to-exploit time.
6. **Targets are marked as targets.** Where the framework sets a standard (for example, that top-scoring protection acts automatically at machine speed), it is a design choice and is labeled so in the statement map. It is not presented as a finding.
7. **Design examples are not measurements.** A vendor tool is cited as an example of a design, never for its performance or timing claims unless a primary source backs them.

2. The source record

2.1 Incident data and government sources

Verizon, 2026 Data Breach Investigations Report, executive summary (data from Oct 2024 to Nov 2025). Type: vendor, based on incident and vulnerability data.
https://www.verizon.com/business/resources/executivebriefs/2026-dbir-executive-summary.pdf

- Exploitation of vulnerabilities is the most common initial access vector in breaches: **31%**. Credential abuse, the previous leader, fell to 13%.
- Only **26%** of critical vulnerabilities (those on CISA's Known Exploited Vulnerabilities list) were fully fixed in 2025, down from 38%.
- The **median time for full resolution rose to 43 days, from 32**. In the median case, organizations had 50% more critical vulnerabilities to patch.
- Third parties were involved in **48%** of breaches, up 60%. Only 23% of third-party organizations fully fixed missing or poorly secured multifactor authentication on cloud accounts, and fixing half of weak-password and permission findings took almost eight months.
- Scope: the 43 days applies to the Known Exploited list, not to every vulnerability.
- Read: executive summary in full. The full report is gated and was not opened.

CERT-In, "AI Agents as Attack Vectors," advisory CIAD-2026-0049, Oct 15, 2026 (rated Critical). Type: government (India's national response team). https://www.cert-in.org.in/s2cMainServlet?pageid=PUBVLNOTES02&VLCODE=CIAD-2026-0049

- Agentic tooling "compressed the time between initial access and final impact from days to a few hours."
- Victims that remediated incompletely "were attacked again within days, at greater scale."
- In incidents disclosed by AI model providers, in which agents exceeded their assigned authorization, none of the affected organizations detected the access. Each was reported by the provider.
- The weaknesses were "conventional": credentials exposed in code and configuration, forgotten test APIs reachable from the internet, plaintext payment-API keys, over-permissive cloud roles, and internet-exposed admin interfaces with weak credentials or no multifactor authentication.
- Advice: default-deny network egress "enforced at the network layer, not through instructions to the model," a short-lived identity for each agent run, and treating an environment as compromised after any credential exposure until forensically cleared.
- Read: in full.

NIST IR 8547, Transition to Post-Quantum Cryptography Standards (Oct 2026, final). Type: standards body. https://doi.org/10.6028/NIST.IR.8547

- Quantum-vulnerable public-key cryptography is to be deprecated after 2030 and disallowed after 2035.
- Related standards: FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA), finalized in 2024.
- It does not say when a quantum computer capable of breaking today's encryption will exist. No source gives a date, and the framework assumes none.

CISA, NSA and NIST, "Quantum-Readiness: Migration to Post-Quantum Cryptography," 2023. Type: government guidance. The government reference for inventorying where public-key cryptography is used, planning the migration, and the concern that encrypted data can be collected now and decrypted later. Not re-read for this paper.

2.2 Vendor research on AI-driven attacks

GreyNoise, "Agents Gone Wild: An AI-Orchestrated Global Campaign Against PaperCut NG/MF," Sept 9, 2026. Type: threat-intelligence vendor. https://www.greynoise.io/blog/ai-orchestrated-campaign-against-papercut-ng-mf

- Hundreds of AI agents (OpenAI's Codex as the harness, a DeepSeek model) compromised at least 440 PaperCut instances at 395 identified organizations in 48 countries.
- The attacker went from an empty workspace to a first remote-code-execution against a real victim in under four hours, and to a first domain administrator in about two more hours.
- Once the full campaign launched, it compromised **at least 11 organizations in 26 seconds**. This is a burst against targets already found and prepared, not a disclosure-to-exploit time.
- In one case it went from initial access to domain administrator in seven minutes, against a high school. It reached domain administrator at only 12 victims. A web application firewall (Cloudflare) defeated its attempt against a protected host.
- Read: primary post in full.

Microsoft Threat Intelligence, "AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise," Oct 21, 2026. Type: vendor, Microsoft's own telemetry. https://www.microsoft.com/en-us/security/blog/2026/10/21/storm-0940-uses-artex-agentic-ai-connectwise-supply-chain-compromise/

- Microsoft assesses that Storm-0940, a China-state aligned actor, used ARTEX, an open-source AI penetration-testing framework, against hundreds of organizations from as early as September 2026.
- In multiple cases the time from initial targeting to privileged access to an organization's cloud was **under three hours**.
- After a supply-chain compromise at a backup vendor (which published its own advisory on Oct 16, 2026), the actor tried to sign in to more than 2,250 tenants using stolen application credentials.
- Microsoft says the speed "allows very little time for Defenders to respond to those alerts." Most of the targeting failed: one password-spray produced 151,000+ failed logons (peaking at 2,900 a minute) and ended in lockouts, with no successful sign-in.
- Read: in full.

Microsoft On the Issues, Mike Yeh, "Security fundamentals to enable the frontier," Nov 19, 2026. Type: vendor commentary. https://blogs.microsoft.com/on-the-issues/2026/11/19/security-fundamentals-to-enable-the-frontier/

- Four moves: address tech debt, close the patching gap, radical observability (treat AI agents as insider risks) and defend at machine speed. Boards should know how many AI agents run in the organization, what they can reach, what they can do on their own, and how they are monitored. These board questions are the basis of the framework's A4 test.
- The post says AI can speed the path from disclosure to exploitation "in minutes" and gives no figure or source. It is not used as a measurement.
- Read: in full.

Israel Defense, Kirill Gelfand (TrendAI), "How Fast Can You Patch a Vulnerability? That's No Longer the Point," Nov 25, 2026. Type: trade press, written by a virtual-patching vendor. https://www.israeldefense.co.il/en/node/73216

- Measure time to protection, not only time to patch. Compensating controls need an owner, an approval process and an expiry date, or they become an excuse to postpone the patch.
- No figures. "Time to protection" is the author's phrase.
- Read: in full.

2.3 Design examples from vendors

Cisco, Jeetu Patel, "Expanding the Defender's Window with Self-defending Infrastructure," Nov 23, 2026 (announcing Shirudo, an open-source project co-developed with OpenAI). Type: vendor. https://blogs.cisco.com/news/expanding-the-defenders-window-with-self-defending-infrastructure

- Used as a design example for S1 and A4. AI agents scan open-source code and vendor disclosures, trace each vulnerability to the affected assets and attack paths, and write a small, targeted shield, built on eBPF, that blocks a path to vulnerable code on a running Linux host without a restart. The agents check that the shield covers the exploit and that the service still runs. The shield "only graduates to active blocking when the evidence says it's safe, and it retires when the real patch lands." Every control is signed, and "proof of where a control came from should travel with it, from the agent that wrote it to the host that runs it." The post names agents as a risk vector that must be "prevented from becoming inadvertent adversaries."
- Caveat: the post's time statements ("minutes and hours," exploit campaigns "within a day of disclosure") carry no primary source, and the project's code and license were not reviewed. We did not test it.
- Read: in full.

Cisco, Steve Clayton, "Moving Fast and Staying in Control in the Agentic AI Era," Nov 9, 2026 (Cisco Live Melbourne). Type: vendor. https://blogs.cisco.com/news/moving-fast-and-staying-in-control-in-the-agentic-ai-era

- Used for two design points: a red-team tool that tests only systems the customer owns and has registered, with every test tied to a named person and stoppable at any time (S2 guardrails); and a playbook that flags whether network equipment's encryption is ready for quantum computing (K2).
- Its statement that attackers "typically strike before a fix is even released" carries no source we have read and is not used.
- Read: in full.

2.4 Patch-cycle and exploit-window data

Rapid7 Labs, "Why Q2 2026 signals the end of traditional patch cycles." Type: vendor. https://rapid7.com/blog/post/tr-new-report-ai-threats-q2-2026-ends-traditional-patch-cycles

- High and critical disclosures (CVSS 7 to 10) doubled, from 4,268 in Q2 2025 to 8,539 in Q2 2026.
- Newly exploited vulnerabilities "held roughly steady (40)." Rapid7's takeaway: "The takeaway isn't that exploitation exploded - it's that disclosure volume is far outstripping what any team can triage."
- 62% of exploited vulnerabilities needed no user interaction, up from 53%.
- SecurityWeek (Kevin Townsend, Aug 18, 2026) quotes Rapid7's Christiaan Beek: "If you still believe we have a monthly patch cycle, forget it."
- Read: Rapid7's post and the SecurityWeek article in full. The full Rapid7 report was not opened.

Mandiant / Google Cloud, Casey Charrier and Robert Weiner, "How Low Can You Go? An Analysis of 2023 Time-to-Exploit Trends," Oct 15, 2024. Type: vendor. https://cloud.google.com/blog/topics/threat-intelligence/time-to-exploit-trends-2023

- Average time to exploit fell from 63 days (2018 to 2019) to 32 days (2021 to 2022) to **5 days in 2023**, across 138 vulnerabilities. The 5-day average leaves out 15 outliers; with them it is 47 days.
- Where the exploit appeared before exploitation, the median was 7 days from disclosure to a public exploit and 30 more to first known exploitation. Where exploitation came first, the median from disclosure to exploitation was 15 days.
- Scope: 2023 data, an average and not a median, before the AI-driven campaigns above.
- Read: primary post in full.

Cloud Security Alliance AI Safety Initiative, "The Collapsing Exploit Window," Apr 11, 2026. Type: industry body. https://labs.cloudsecurityalliance.org/research/csa-whitepaper-collapsing-exploit-window-ai-mtte-20260411-cs/

- Argues the grace period collapsed from the roughly 32-day patch median to about 5 days of median time to exploit in 2025. Its 2025 figure rests on a vendor blog. Context only; no figure from it is used.
- Read: in full.

Security Boulevard (Dhingra, Nov 3, 2026, syndicated from Kratikal) and Barracuda (Vizard, updated Nov 19, 2026). Type: vendor blogs. Both say AI tools speed attacks and that teams must weigh outage risk against exposure. Context only.

2.5 Frameworks compared

- **NIST CSF 2.0** (CSWP 29, 32 pages, read in full) and NIST's online Implementation Examples (downloaded and searched). https://doi.org/10.6028/NIST.CSWP.29

- **CIS Critical Security Controls v8.1** (public control list read; the Safeguards were not read). https://www.cisecurity.org/controls/cis-controls-list
- **MITRE ATT&CK, D3FEND, Engage** (Engage v1.0 data read from MITRE's published data files). https://attack.mitre.org · https://d3fend.mitre.org · https://engage.mitre.org
- The scoring and search record are in the comparison scoring paper.

3. What the evidence shows, and its limits

What it shows

- Exploiting vulnerabilities is now the leading way into a breach (Verizon: 31%), and full fixes are getting slower (26% fully fixed; 43 days median).
- AI-driven campaigns run on a clock of hours (GreyNoise: under 4 hours to a first intrusion; Microsoft: under 3 hours to privileged cloud access; CERT-In: "a few hours").
- The attacker and defender clocks differ by an order of magnitude or more: hours against a 43-day median.
- The weaknesses attackers use are conventional (CERT-In), so having a control on a list is not the same as keeping it working.
- Third parties are a large share of the problem (Verizon: 48% of breaches; Microsoft: a compromised backup vendor's application opened its customers).
- The post-quantum standards are final, with dates for retiring today's algorithms (NIST IR 8547).
- A major vendor and an AI lab are building agent-written, eBPF-based protection with observe-then-enforce, expiry and signed provenance (Cisco and OpenAI). This supports the S1 design as a direction, not as a result.

Limits

- Most sources are vendors, who sell security products.
- The 26-second figure measures a burst against prepared targets. There is no sourced disclosure-to-exploit time for AI-driven campaigns.
- Most attacks in the GreyNoise and Microsoft accounts failed against lockouts and a web application firewall.
- Rapid7 counts about the same number of newly exploited vulnerabilities as the year before, while disclosures doubled.
- No source gives a date for a quantum computer that breaks today's encryption. The framework argues from how long data must stay secret and how long migration takes.
- The S1 and A4 requirements about signing and observe-then-enforce have not been tested in practice.

4. Statement-to-source map

Statement in the framework paper	Source
43 days median to fully fix critical flaws; 26% fully fixed; 31% exploitation	Verizon 2026 DBIR
Third parties in 48% of breaches	Verizon 2026 DBIR
Under 4 hours to first intrusion; 11 organizations in 26 seconds	GreyNoise, Sept 9, 2026
Under 3 hours to privileged cloud access; 2,250+ tenants tried	Microsoft Threat Intelligence, Oct 21, 2026

Statement in the framework paper	Source
Days to “a few hours”; conventional weaknesses; default-deny egress; short-lived identity per agent run	CERT-In CIAD-2026-0049
Disclosures doubled; exploited count steady at 40	Rapid7 Labs, Q2 2026
Average time to exploit 5 days in 2023 (47 with outliers)	Mandiant, Oct 2024
“Time to protection”; expiry dates on compensating controls	Gelfand, Israel Defense, Nov 25, 2026
Board questions on AI agents	Microsoft On the Issues, Nov 19, 2026
Deprecate after 2030, disallow after 2035; FIPS 203, 204, 205	NIST IR 8547; NIST FIPS
Collect now, decrypt later; inventory and migration planning	CISA, NSA and NIST (2023)
NIST outcomes “are not a checklist of actions to perform”; no scores, no response times, no quantum	NIST CSWP 29 and online Implementation Examples (comparison scoring paper)
CIS Control 18 is penetration testing	CIS Controls v8.1
Shield starts in observe mode, enforces on evidence, retires at patch; controls signed with provenance	Design example: Cisco and OpenAI, Nov 23, 2026. The requirement itself is the IQ4hire standard
Red-team tests only on owned, authorized systems, tied to a named person, stoppable	Design example: Cisco, Nov 9, 2026; CERT-In advice
Top-scoring protection is pre-approved and acts automatically at machine speed	IQ4hire standard (a target, not a finding). Reviewed as attacker data changes
Exposure test for long-lived secrets: shelf life plus migration time against time to a capable quantum computer	IQ4hire reasoning. No date for a quantum computer is claimed
Toolbox categories and open standards (framework paper, Appendix C)	Public standards and projects, named only; not tested or endorsed

5. Sources not read in full

The full Verizon and Rapid7 reports and the CIS Safeguards were not opened. The CISA, NSA and NIST quantum-readiness guidance is cited from earlier reading and was not re-read for this paper. The code and license of the Cisco and OpenAI project were not reviewed.

6. References

1. NIST. (2024). The NIST Cybersecurity Framework (CSF) 2.0. NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29
2. NIST. CSF 2.0 Reference Tool: Implementation Examples and Informative References. https://csrc.nist.gov/projects/cybersecurity-framework/filters
3. Center for Internet Security. CIS Critical Security Controls v8.1. https://www.cisecurity.org/controls/cis-controls-list
4. MITRE. ATT&CK. https://attack.mitre.org · D3FEND. https://d3fend.mitre.org · Engage. https://engage.mitre.org
5. NIST. (2024). FIPS 203, Module-Lattice-Based Key-Encapsulation Mechanism Standard. https://csrc.nist.gov/pubs/fips/203/final
6. NIST. (2024). FIPS 204, Module-Lattice-Based Digital Signature Standard. https://csrc.nist.gov/pubs/fips/204/final

7. NIST. (2024). FIPS 205, Stateless Hash-Based Digital Signature Standard. https://csrc.nist.gov/pubs/fips/205/final
8. Moody, D., Perlner, R., Regenscheid, A., Robinson, A. and Cooper, D. (October 2026). Transition to Post-Quantum Cryptography Standards. NIST IR 8547. https://doi.org/10.6028/NIST.IR.8547
9. CISA, NSA and NIST. (2023). Quantum-Readiness: Migration to Post-Quantum Cryptography. https://www.cisa.gov/sites/default/files/2023-08/Quantum%20Readiness_Final_CLEAR_508c%20%283%29.pdf
10. Verizon. (2026). 2026 Data Breach Investigations Report: Executive Summary. https://www.verizon.com/business/resources/executivebriefs/2026-dbir-executive-summary.pdf
11. GreyNoise. (September 9, 2026). Agents Gone Wild: An AI-Orchestrated Global Campaign Against PaperCut NG/MF. https://www.greynoise.io/blog/ai-orchestrated-campaign-against-papercut-ng-mf
12. CERT-In. (October 15, 2026). AI Agents as Attack Vectors. Advisory CIAD-2026-0049. https://www.cert-in.org.in/s2cMainServlet?pageid=PUBVLNOTES02&VLCODE=CIAD-2026-0049
13. Microsoft Security Research and Microsoft Threat Intelligence. (October 21, 2026). AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise. https://www.microsoft.com/en-us/security/blog/2026/10/21/storm-0940-uses-artex-agentic-ai-connectwise-supply-chain-compromise/
14. Yeh, M. (November 19, 2026). Security fundamentals to enable the frontier. Microsoft On the Issues. https://blogs.microsoft.com/on-the-issues/2026/11/19/security-fundamentals-to-enable-the-frontier/
15. Patel, J. (November 23, 2026). Expanding the Defender's Window with Self-defending Infrastructure. Cisco Blogs. https://blogs.cisco.com/news/expanding-the-defenders-window-with-self-defending-infrastructure
16. Clayton, S. (November 9, 2026). Moving Fast and Staying in Control in the Agentic AI Era. Cisco Blogs. https://blogs.cisco.com/news/moving-fast-and-staying-in-control-in-the-agentic-ai-era
17. Gelfand, K. (November 25, 2026). How Fast Can You Patch a Vulnerability? That's No Longer the Point. Israel Defense. https://www.israeldefense.co.il/en/node/73216
18. Rapid7 Labs. (2026). Why Q2 2026 signals the end of traditional patch cycles. https://rapid7.com/blog/post/tr-new-report-ai-threats-q2-2026-ends-traditional-patch-cycles
19. Townsend, K. (August 18, 2026). AI-Driven Vulnerability Surge Breaks the Traditional Patching Model. SecurityWeek. https://www.securityweek.com/ai-driven-vulnerability-surge-breaks-the-traditional-patching-model/
20. Charrier, C. and Weiner, R. (October 15, 2024). How Low Can You Go? An Analysis of 2023 Time-to-Exploit Trends. Mandiant / Google Cloud. https://cloud.google.com/blog/topics/threat-intelligence/time-to-exploit-trends-2023
21. Cloud Security Alliance AI Safety Initiative. (April 11, 2026). The Collapsing Exploit Window. https://labs.cloudsecurityalliance.org/research/csa-whitepaper-collapsing-exploit-window-ai-mtte-20260411-cs/
22. Dhingra, S. (November 3, 2026). AI Can Now Hack Faster Than Your Team Can Patch. Security Boulevard. https://securityboulevard.com/2026/11/ai-can-now-hack-faster-than-your-team-can-patch-heres-what-to-do-about-it/
23. Vizard, M. (updated November 19, 2026). Patch management vs. IT outages: balancing cyber risk. Barracuda Networks Blog. https://blog.barracuda.com/2026/11/19/patch-management-it-outages-ai-era

Notices

IQ4hire is not affiliated with, endorsed by or sponsored by NIST, the Center for Internet Security, MITRE, Cisco, OpenAI, Microsoft, Verizon, GreyNoise, CERT-In or any other organization named in this paper. Product, framework and company names belong to their owners and are used only to identify the works discussed. Quotations are taken from the published documents, with sources cited. The scoring in this paper is IQ4hire's own method and opinion, and it is open to anyone to repeat or dispute.