

PERSPECTIVE · OCTOBER 2026

Health Insurers and *Cyber Hygiene*

Prime targets, basic entry routes and the case for AI-equipped White Box Purple Testing

Lessons from 29 publicly reported breaches, 2010 to 2026

Executive Summary and Recommendations

- **Health insurers sit on a mine of pure gold.** If data is gold, a health insurer holds the purest seam: identity, bank details and medical history, linked together and kept for years.
- **They are prime targets, and the industry knows it.** Change Healthcare (192.7 million people), Medibank (9.7 million) and DentaQuest (at least 15 million) show the scale. In June 2025 Google warned that Scattered Spider was targeting US insurers.
- **Most attackers logged in rather than broke in.** Of the 16 intrusions whose entry route is public, 9 (56%) began with phishing, stolen or socially engineered credentials, or a default password. Remote access protected by a password alone recurs.
- **Slow detection and supplier exposure made the losses larger.** Attackers went unnoticed for months and sometimes years, and several of the largest breaches occurred at claims processors and benefits administrators acting for insurers.
- **AI is shortening the time defenders have.** Critical vulnerability disclosures have almost doubled to nearly 20 a day since April 2026, and published estimates of time-to-exploit now range from hours to below zero. Patching discipline joins identity and monitoring as a core hygiene control.
- **We recommend White Box Purple Testing.** Red and blue teams work together with full knowledge of the estate, both equipped with frontier or near-frontier AI, so that the insurer finds and fixes weaknesses at the speed an AI-assisted attacker will look for them. The method builds on the purple-team exercise that TIBER-EU already requires and adds white-box scope, AI on both sides and metrics tied to the failures in this record.

Health insurers hold the most complete data an attacker can steal

Most organisations hold one kind of valuable data: a bank holds financial data and a hospital holds clinical records. A health insurer holds both and links them to a verified identity: name, date of birth, address, government identifiers, bank details, claims history, diagnoses and

prescriptions. It keeps them for members, former members, dependants and employees, often for many years.

The combined record is worth more to a criminal than any of its parts. A stolen card can be cancelled within hours; an exposed diagnosis stays exposed. Records of this kind support identity theft and insurance fraud and, as Medibank's customers found in 2022, public release of the most sensitive claims when a ransom is refused.

Attackers know this, and the record shows it. The largest health data breach in US history, the 2024 attack on Change Healthcare, a claims clearinghouse owned by UnitedHealth Group, affected 192.7 million people. In Australia, the 2022 Medibank breach affected 9.7 million current and former customers. In France, two third-party payment operators serving health insurers exposed 33 million people within a week in 2024. After law enforcement disrupted the ALPHV ransomware operation in December 2023, its administrator urged affiliates to target hospitals, and healthcare became its most victimised sector. In June 2025 Google warned the US insurance industry that intrusions bearing the hallmarks of Scattered Spider had reached insurers; Aflac, Erie Insurance and Philadelphia Insurance reported intrusions that month, and Aflac described its own as part of a cybercrime campaign against the insurance industry.

Most attackers logged in rather than broke in

We compiled and verified a chronology of 29 publicly reported cyber incidents at health insurers, health plans and the processors acting for them, between 2010 and 2026, across the United States, the United Kingdom, Australia, France, India and Canada. The Appendix lists them with our coding rules. Two of the 29 (Kaiser and Blue Shield of California) are privacy disclosures through website tracking code and do not count as intrusions, so we leave them out of the entry-route count. Of the remaining 27, the public record states how the attacker got in for 16.

ENTRY ROUTE	INCIDENTS	SHARE	EXAMPLES
Phishing, stolen or socially engineered credentials, or default password	9	56%	Anthem, Premera, Bankers Life, Utah Medicaid, Medibank, Viamedis, Change Healthcare, Aflac, DentaQuest
Vulnerable or misconfigured web application	4	25%	WellPoint, Aetna, Religare, guard.me
Insider	3	19%	CalOptima, Bupa Global, Humana subcontractor

Three patterns recur across these cases:

Remote access protected by a password alone. Change Healthcare's attackers used stolen credentials on a remote-access portal that had no MFA, as UnitedHealth's chief executive told Congress. The Australian Information Commissioner alleges that Medibank allowed remote access without MFA, and that a contractor's stolen credentials opened an administrator account with reach across most of its systems. US agencies report the same pattern for the Daixin and Qilin ransomware groups, both active against the health sector.

Social engineering of staff and help desks. Where MFA was in place, attackers persuaded people to give it away. Aflac's announcement states that the attacker used social engineering tactics to gain access to its network. DentaQuest states that one employee handed over credentials and an MFA code. Insurers run large help desks and call centres, and groups such as Scattered Spider and ShinyHunters specialise in exploiting them.

Patching rarely appears in the disclosed record. None of the 16 is traced on the public record to an unpatched known vulnerability at the insurer itself; the clearest patching event, the 2023 MOVEit campaign, was a zero-day that reached health plans through vendors such as Welltok (8.5 million people). We read this as a feature of what insurers and regulators choose to disclose rather than as evidence that patching is sound, since the entry route is not public for 11 of the 27 intrusions. Across all sectors, Mandiant's investigations for 2025 found exploits to be the most common entry route, at 32% of intrusions, followed by voice phishing at 11%.

Slow detection and supplier exposure made the losses larger

Attackers went unnoticed for long periods: about nine months at Premiera, about 20 months at Excellus, up to nine years at Dominion National and about ten months at a TriZetto eligibility portal. In its settlements with Anthem, Premiera and Excellus, the US Department of Health and Human Services (HHS) found the same root causes in each case: no enterprise-wide risk analysis, weak risk management and inadequate review of system activity. The UK Information Commissioner found that Bupa did not monitor bulk extraction from its customer system, which allowed an employee to copy the records of 547,000 customers.

The largest exposures sat with suppliers. Change Healthcare processes claims for much of the US market; Episource (6.7 million people) codes records for health plans; MedImpact manages pharmacy benefits; DentaQuest administers dental benefits for Medicaid and Medicare; Viamedis and Almerys process third-party payments for French insurers. In each case members were exposed through a firm whose controls the insurer did not operate but for which it remained accountable.

AI is shortening the time defenders have

The 2026 evidence from public authorities, AI laboratories and incident responders is consistent in direction, if not in every number:

- **Discovery:** in Anthropic's Project Glasswing, the Mythos Preview model estimated 6,202 high or critical vulnerabilities across more than 1,000 open-source projects by May 2026. Independent firms confirmed 62.4% of the 1,752 they assessed as high or critical. Of 530 disclosed to maintainers, 75 had been patched. The constraint is shifting from finding flaws to fixing them.
- **Volume:** the Bank for International Settlements reports critical vulnerability disclosures rising from about 7 a day (2018 to 2021) and 10 a day (2022 to 2025) to almost 20 a day after 1 April 2026. The UK National Cyber Security Centre has told organisations to prepare for a "vulnerability patch wave".
- **Speed:** estimates of mean time-to-exploit vary with method. The South African Reserve Bank reports a fall from 2.3 years in 2018 to about 20 hours by early 2026; Mandiant estimates minus seven days for 2025, meaning exploitation routinely begins before a patch is released.
- **Capability:** the UK AI Security Institute estimates that the length of cyber tasks AI models can complete with 80% reliability doubled every 4.7 months from late 2024, and that Claude Mythos Preview and GPT-5.5 substantially exceeded that trend.

This has two consequences for health insurers. First, patching discipline becomes a hygiene control on the same footing as MFA, because an exposed system that waits weeks for a patch can be found by an AI-assisted attacker within hours. Second, the social-engineering routes that already dominate the record become cheaper to run at scale. The Financial Stability Board's June 2026 consultation on AI sound practices includes "phishing-resistant multi-factor authentication (MFA) to mitigate some of the risks of advanced AI models" among its cyber controls. Controls calibrated to human-paced attackers leave less margin against AI-paced ones.

White Box Purple Testing lets insurers find and fix weaknesses at AI speed

Financial-sector schemes such as CBEST in the UK, and TIBER-EU and DORA threat-led penetration testing in the EU, test whether a stealthy, intelligence-led attacker can reach critical functions. They are valuable, and TIBER-EU already makes a purple-teaming exercise a fixed element of its closure phase. They are, however, periodic and designed around a human-paced adversary.

We recommend adding what we call White Box Purple Testing. Red and blue teams already collaborate in purple exercises; the distinctive elements here are white-box scope, AI on both sides and metrics tied to the failures in this record:

- **White box:** the testers receive full knowledge of the estate, including architecture, identity and access configuration, supplier connections, code and logs. The aim is coverage of the whole estate rather than a test of stealth.
- **Purple:** red and blue teams work side by side. Every attack path the red team finds is checked against what the blue team detected and fixed while the test is running.
- **AI on both sides:** both teams use frontier or near-frontier AI. The red team uses it to find and chain weaknesses at machine speed; the blue team uses it to triage alerts, write detections and prioritise fixes. The exercise measures whether the insurer can close a gap faster than an AI-assisted attacker can find it.

Three design principles keep the exercise safe and useful:

- **Contain the AI.** In July 2026 the UK AI Security Institute reported that AI agents under cyber testing took unsanctioned real-world actions in 10 of 122 runs, including an attempt to insert malicious code into an open-source project. AI-equipped testing should run in isolated ranges or under strict egress control, with real-time monitoring of agent actions and a human control team with authority to stop it.
- **Protect member data.** Use synthetic or de-identified data, or models deployed within the insurer's own environment, so that health data stays under the insurer's control throughout the test.
- **Measure what the record shows matters.** Track the share of logins and supplier connections on phishing-resistant MFA, whether the help desk resists a reset request made by telephone, the time to detect bulk data extraction, the time to patch internet-facing systems and the share of the estate reviewed with AI each quarter.

Questions for the Board this quarter

- Is every remote-access route, including supplier and contractor access, protected by phishing-resistant MFA?
- Could someone reset an executive's MFA by telephoning our help desk, and when did we last test that?
- How long would an intruder copying member data go unnoticed, and what evidence supports that answer?
- Which suppliers would stop claims or expose most member data if breached, and do our contracts require MFA, logging and prompt notification?

- How many days do we take today to patch an internet-facing system, and what would an exploitation window of hours require?
- When did our red and blue teams last work together with AI tools, and what did they fix?

Health insurers will always hold valuable data. The controls that make it expensive to steal are well understood; the record shows that they need to be applied consistently, tested regularly and, from now on, tested at AI speed.

We will be happy to bring the benefit of our experience and regulatory expertise to health insurers, their suppliers and their boards, including the design and running of White Box Purple Testing.

Santosh Pandit · Board and CxO Advisor · Founder, BeatQuantum Limited

BeatQuantum.com | SantoshPandit.com | contact@santoshpandit.com

Appendix: the 29 incidents and how we coded them

Coding rules: an entry route is coded only where the company, a regulator, a court filing, a government advisory or sworn testimony states it, with one exception noted below. “Not public” means no such statement exists. Tracking-code disclosures are privacy incidents rather than intrusions and are excluded from the entry-route count. Insider means an employee or subcontractor misused authorised access. Full details and sources for each incident are in our breach chronology, available on request.

#	YEAR	ORGANISATION	ENTRY ROUTE AS CODED
1	2010	WellPoint (USA)	Web application: data exposed after a software upgrade
2	2012	Utah Department of Health, Medicaid (USA)	Credentials: default password on a server
3	2015	Anthem (USA)	Credentials: spear-phishing
4	2015	Premera Blue Cross (USA)	Credentials: phishing email installing malware
5	2015	CareFirst BlueCross BlueShield (USA)	Not public
6	2015	Excellus BlueCross BlueShield (USA)	Not public
7	2016	Newkirk Products, ID-card vendor (USA)	Not public
8	2016	SCAN Health Plan (USA)	Not public
9	2016	CalOptima (USA)	Insider
10	2017	Aetna (USA)	Web application: documents served without login
11	2017	Bupa Global (UK)	Insider
12	2018	Bankers Life (USA)	Credentials: compromised employee credentials
13	2019	Dominion National (USA)	Not public
14	2020	Religare Health Insurance (India)	Web application: misconfigured server (threat-intelligence source, the one exception)
15	2021	guard.me (Canada)	Web application: website vulnerability
16	2021	Humana, via subcontractor (USA)	Insider
17	2022	Medibank (Australia)	Credentials: stolen contractor login, VPN without MFA
18	2023	Harvard Pilgrim Health Care (USA)	Not public
19	2023	MCNA Dental (USA)	Not public
20	2024	Viamedis and Almerys (France)	Credentials: phished health-professional logins
21	2024	Change Healthcare (USA)	Credentials: stolen, remote-access portal without MFA
22	2024	Kaiser Foundation Health Plan (USA)	Excluded: tracking code
23	2024	Star Health (India)	Not public (insider involvement alleged)
24	2025	Blue Shield of California (USA)	Excluded: tracking code
25	2025	Episource (USA)	Not public

#	YEAR	ORGANISATION	ENTRY ROUTE AS CODED
26	2025	Aflac (USA)	Credentials: social engineering
27	2025	MedImpact (USA)	Not public
28	2026	TriZetto Provider Solutions (USA)	Not public
29	2026	DentaQuest (USA)	Credentials: employee socially engineered for login and MFA code

Totals: credentials 9, web application 4, insider 3, not public 11, excluded 2.

References

1. HHS Office for Civil Rights, breach portal (US incident counts).
https://ocrportal.hhs.gov/ocr/breach/breach_report.jsf
2. Change Healthcare victim count rises to 192.7 million, TechTarget.
<https://www.techtarget.com/healthtechsecurity/news/366628716/Change-Healthcare-data-breach-victim-count-rises-to-193-million>
3. Change Healthcare breached via Citrix portal with no MFA, TechTarget.
<https://www.techtarget.com/searchsecurity/news/366582824/Change-Healthcare-breached-via-Citrix-portal-with-no-MFA>
4. FBI, CISA and HHS, #StopRansomware: ALPHV Blackcat, AA23-353A (updated February 2024).
<https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-353a>
5. Information Commissioner v Medibank, concise statement summary, JWS. <https://jws.com.au/what-we-think/aic-v-medibank-concise-statement-released-oaic/>
6. Scattered Spider targeting insurance sector, The Record. <https://therecord.media/scattered-spider-targeting-insurance-sector-following-retail-attacks>
7. Aflac Incorporated Discloses Cybersecurity Incident, Aflac, 20 June 2025.
<https://newsroom.aflac.com/2025-06-20-Aflac-Incorporated-Discloses-Cybersecurity-Incident>
8. DentaQuest data theft hack affects 15 million, GovInfoSecurity.
<https://www.govinfosecurity.com/dentaquest-data-theft-hack-affects-15m-patients-a-32390>
9. Viamedis and Almerys: 33 million French citizens affected, Dark Reading.
<https://www.darkreading.com/cloud-security/33m-french-citizens-countrys-largest-ever-breach>
10. Anthem record USD 16 million HIPAA settlement and OCR findings, GovInfoSecurity.
<https://www.govinfosecurity.com/anthem-mega-breach-record-16-million-hipaa-settlement-a-11622>
11. Premera Blue Cross pays USD 6.85 million, HHS. <https://www.hhs.gov/guidance/document/health-insurer-pays-685-million-settle-data-breach-affecting-over-104-million-people>
12. Excellus will pay USD 5.1 million to OCR, Healthcare IT News.
<https://www.healthcareitnews.com/news/excellus-will-pay-51m-ocr-after-data-breach-affects-93m-people>
13. Dominion National nine-year breach, HIPAA Journal. <https://www.hipaajournal.com/members-affected-by-dominion-national-9-year-phi-breach/>
14. TriZetto Provider Solutions breach, Security Affairs. <https://securityaffairs.com/189149/data-breach/cognizants-trizetto-provider-solutions-data-breach-impacted-over-3-4-million-patients.html>
15. Bupa fined by ICO after employee stole customer information, Digital Health.
<https://www.digitalhealth.net/2018/10/bupa-fine-ico-employee-stole-information/>
16. Episource data breach, HIPAA Journal. <https://www.hipaajournal.com/episource-data-breach/>

17. Welltok MOVEit data breach, HIPAA Journal. <https://www.hipaajournal.com/welltok-data-breach/>
18. Project Glasswing initial update, Anthropic, May 2026. <https://www.anthropic.com/research/glasswing-initial-update>
19. A Mythos moment? Frontier AI and cyber risk, BIS Bulletin No 129, 20 July 2026. <https://www.bis.org/publications/bulletin-129-mythos-moment-frontier-ai-and-cyber-risk.pdf>
20. Prepare for the vulnerability patch wave, NCSC. <https://www.ncsc.gov.uk/blogs/prepare-for-vulnerability-patch-wave>
21. Financial Stability Review, June 2026, South African Reserve Bank. <https://www.resbank.co.za/content/dam/sarb/publications/reviews/finstab-review/2026/first-edition/first-fsr-june.pdf>
22. How fast is autonomous AI cyber capability advancing?, UK AI Security Institute, 13 May 2026. <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>
23. Sound practices for responsible adoption of AI, consultation report, FSB, June 2026. <https://www.fsb.org/uploads/P100626.pdf>
24. TIBER-EU Purple Teaming Guidance, ECB, January 2025. https://www.ecb.europa.eu/pub/pdf/annex/ecb.tiber_eu_purple_best_practices_2025.en.pdf
25. Incident report: unsanctioned agent behaviour during cyber testing, UK AI Security Institute. <https://www.aisi.gov.uk/blog/incident-report-unsanctioned-agent-behaviour-during-cyber-testing>
26. M-Trends 2026, Mandiant (Google Cloud). <https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026>