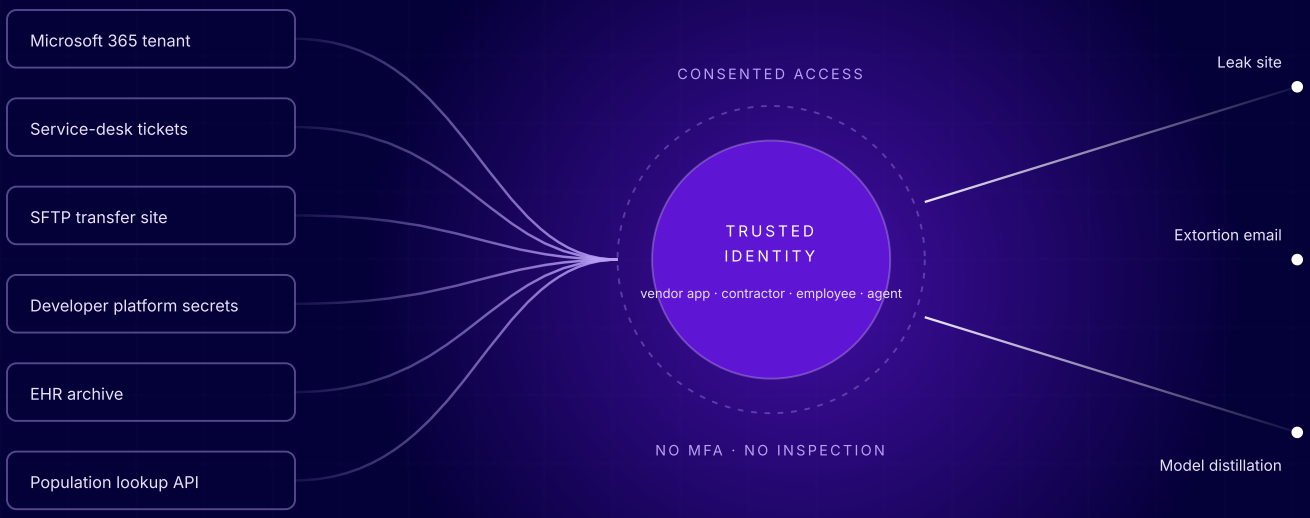
Nightfall AI

BREACH REPORT

SEPTEMBER + OCTOBER 2026

Nobody broke in. *They were let in.*

Twenty-eight breaches, five patterns, one missing control. What two months of vendor tokens, stolen secrets, insider reads and AI agents taught us about where enterprise data actually leaves.



28
incidents in the ledger

1 in 3+
came through a trusted third party's access

60M+
people in confirmed exposures

AI moves your data. Nightfall controls it.

NIGHTFALL.AI

01 · EXECUTIVE SUMMARY

The vendor is the *new perimeter.*

Most of the breaches that mattered in September and October 2026 were not failures of perimeter defense. A legitimate identity, often belonging to a vendor, a contractor, or a non-human app, was used to pull data out of a system that had been doing its job quietly for years.

28

incidents in our ledger, disclosed or materially updated between Sept 1 and Oct 31

10 of 28

reached the victim through a third party's platform or credential, the largest confirmed vector

60M+

people in confirmed exposure counts, before attacker-claimed figures

Source: Nightfall analysis of company filings, regulator notices and primary reporting. Methodology on page 3.

Five patterns stood out

The vendor is the perimeter. ConnectWise, Veradigm, AdaptHealth, iRhythm, EY, Oracle Health and California's ETS were all breached through a third party's platform or credential.

Credentials and secrets opened most doors. Railway, Zeabur and JetBrains exposed developer secrets at scale. Denmark's population register was emptied with a former employee's password.

AI is on both sides. Attacker-side agents (ARTEX, Codex swarms) compromised hundreds of organizations in days. Employee-side AI misuse caused the first regulator-logged breaches of its kind.

Insiders are a data problem. Fifteen NHS staff were dismissed for reading the records of the 2023 Nottingham attack victims. Nothing was exfiltrated; the records were simply reachable.

Regulated data sat unwatched. IDScan's cloud platform was open five months. The California student-data theft ran nearly five. Oracle Health's 19.9 million count surfaced eighteen months after the breach.

What this paper covers

A sourced ledger of the 28 incidents, the five trends with the evidence behind each, the controls that would have changed outcomes, and how Nightfall maps to them.

> *Detection after the fact did nothing for the organizations in this report. The value was always in the moment before the data left.*

IN THIS PAPER

02 The breach ledger

03 Patterns by the numbers

04 Trend 1: The vendor is the perimeter

05 Trend 2: Secrets are data

06 Trend 3: Valid identity, wrong records

07 Trend 4: AI on both sides

08 Trend 5: Regulated data, unwatched

09 What would have changed the outcome

10 Where Nightfall fits

11 Sources

02 · THE BREACH LEDGER

Twenty-eight incidents, *one shape.*

Newest first. Scope: incidents disclosed or materially updated Sept 1 to Oct 31, 2026, involving sensitive data in SaaS, cloud workspaces, developer environments, AI tools or insider access. Figures come from filings, regulator notices, vendor disclosures or original reporting that cites them; attacker-only claims are labeled.

DISCLOSED	ORGANIZATION	HOW THE DATA LEFT	WHAT WAS EXPOSED	SCALE
Oct 30	**Chainalysis**	Attacker used "offensive agentic techniques" to reach internal tools; Nov 4 update confirmed password hashes and live tokens for all users	Names, usernames, emails, role labels, hashes, tokens	Not disclosed
Oct 29	**Florida Dept. of Education**	Financial-aid website breach; vector undisclosed	SSNs, names, DOBs, addresses	~4.9M
Oct 26	**90+ South Korean orgs**	Former pentester ran the ARTEX autonomous AI framework against hospitals, schools, corporations	Names, resident registration numbers, phones	8.2M records; 38 ransom demands
Oct 22	**Quest Diagnostics**	Unauthorized access to Quanum Lab Services Manager, the client portal for ordering tests and viewing results	Patient records; PHI/PII confirmed, fields undetermined	Not disclosed
Oct 21	**Railway**	Intruder in build infrastructure decrypted customer environment variables; targeted AI provider API keys	API keys, DB passwords, private keys (wallets drained)	Up to 1.85% of workspaces, ~100,000
Oct 20	**Zeabur**	Plaintext long-lived admin credentials; management interfaces on the public internet; second breach in two months	3,218,965 env-var records, 566 plaintext API keys	321,810 accounts; shutting down
Oct 16	**ConnectWise Cloud Backup**	Stolen vendor app credential used to sign in to customer Microsoft 365 tenants as the app; apps granted themselves Exchange Admin and Global Admin	Graph tokens over Exchange, OneDrive, SharePoint, Teams	~1,000 tenants touched; 300+ successful
Oct 16	**California Dept. of Education / ETS**	Compromised login to contractor's SFTP site; copying ran May 22 to Sept 18	Names, DOBs, addresses, IDs, disability status, test data	5.7M+ students
Oct 12	**Denmark CPR / Pays ApS**	Vendor's lawful lookup access abused with a former employee's leaked password; ~14M automated searches in 21 days	Names, addresses, CPR numbers	~8.8M people
Oct 9	**iRhythm**	Social engineering gave access to third-party hosted business apps (June 8); extortion the next day	Names, contact, insurance and device numbers, dates of service	8,179,772
Oct 9	**Holland & Knight, Squire Patton Boggs, Nelson Mullins**	Disclosures filed the same week; Nelson Mullins cited access to a user's laptop	Client data incl. SSNs	Not disclosed
Oct 8	**ASOS**	Attacker posed as a trusted contact, got an employee's login, used it on third-party platforms; threats sent as in-app push notifications	Names, contact details	Not disclosed
Oct 7	**Arizona court system**	Court employee clicked a phishing link (Sept 24)	PII; 150,000+ foster-care review records	~1.3M
Oct 7	**EY**	Third-party IT service management platform; support-ticket attachments held client tax files (Mar 28 to Apr 12)	Tax IDs and financial details of Goldman Sachs wealth and Man Group clients	Not disclosed
Oct 4	**Nikkei**	Google Workspace account accessed from late July; Microsoft 365 account used Sept 30 to send ~9,000 phishing emails	Names, emails of 1,646 contacts; email content	1,646 + recipients

02 · THE BREACH LEDGER, CONTINUED

DISCLOSED	ORGANIZATION	HOW THE DATA LEFT	WHAT WAS EXPOSED	SCALE
Oct 2	**Oracle Health (legacy Cerner)**	Compromised customer credentials used to reach legacy Cerner servers not yet migrated to Oracle's cloud (Jan to Apr 2025); count surfaced via Texas AG portal	Names, SSNs, diagnoses, medications, lab results	19,929,149
Sept 30	**Bee Cheng Hiang**	Employee used a personal generative AI tool to write a mass-mail script; a missing bracket put up to 999 recipients in one To field	Customer email addresses	95,364
Sept 29	**ESNEFT (NHS)**	Staff allegedly read one patient's records without cause; 10 suspended	Medical records	1 patient
Sept 22	**BigCommerce merchants via Ribon**	Third-party app's API credentials compromised through a breach at its parent; scripts injected into storefronts	Customer names, emails, phones, addresses	App on "hundreds" of stores
Sept 16	**Spain AEPD, first AI-agent breach**	Autonomous AI agent scanned, logged in, found a flaw, altered personal records and read invoices (per the victim's notification)	Personal records, invoice data	Not disclosed
Sept 12	**Revolut**	Fraudulent data requests sent from a legitimate government email domain; Revolut complied	Passport and license copies, selfies, statements	"Limited"; high-net-worth users
Sept 10	**PaperCut campaign, 395 orgs**	Hundreds of AI agents (Codex, DeepSeek) built and ran exploits; 11 orgs fell in 26 seconds	Credentials from 280 victims; domain secrets from 147	440 instances, 48 countries
Sept 9	**AdaptHealth**	Social engineering of a third-party contractor's privileged account; access to cloud patient management, document storage and EHR portals	Names, contact, demographics, insurance, health data	4,115,802
Sept 8	**Veradigm**	Credentials stolen from a vendor's network used against a customer-service API	Patient PII incl. some SSNs; no clinical data	Attacker claims 3.5M; company: "small number of customers"
Sept 5	**JetBrains Cadence**	Unpatched TeamCity RCE (CVE-2026-63077, on CISA KEV since Aug 5); intrusion Aug 8 to 24	AWS IAM credentials, S3 files, a full 2024 backup, synced source	Not disclosed
Sept 4	**IDScan.net**	Unauthorized access to the VeriScan cloud platform from April 4 to Sept 2	Names, DOBs, license, passport and ID numbers	Seller claims 153M license scans
Sept 2	**Dropbox via Lenovo ID**	Lenovo email-verification flaw let attackers mint a Lenovo ID for a victim's email and sign in to Dropbox with no password (Aug 4 to 21)	Files viewed and downloaded	~5,000 accounts
Sept 1	**Aesto Health**	Unauthorized access to portions of its AWS environment (Dec 2 to 18, 2025); six months to public notice	Names, SSNs, license numbers, financial accounts, medical and insurance data	9,540,683

Treated separately: Binance, Oct 29

About $6.34 billion. The chain ran through a wallet engineer's tampered Homebrew install, a fake-recruiter repository on a second engineer's machine, an internal AI tool used to pivot to two more machines, and a container deployment token found on one of them. The FBI attributed it to North Korea's TraderTraitor cluster. Its lesson is about developer machines and secrets, so it sits in Trend 2.

Treated separately: Nottingham University Hospitals, Oct

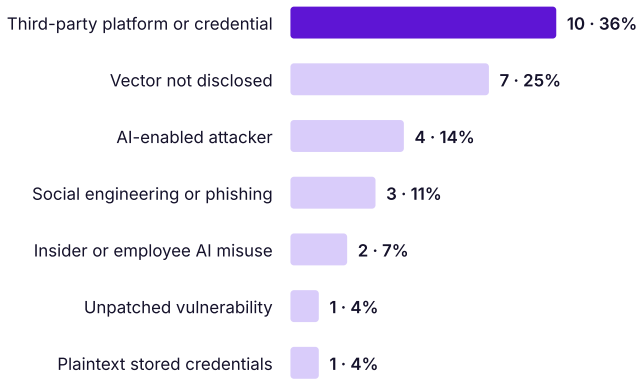
Fifteen staff, including doctors and nurses, dismissed for opening the medical records of the June 2023 Nottingham attack victims with no clinical reason. No external attacker, no exfiltration. It anchors Trend 3.

03 · PATTERNS BY THE NUMBERS

Access was handed over, *not broken.*

Where the vector was disclosed at all, a human giving up access or a machine identity being abused outnumbered software exploitation more than seven to one.

Primary initial vector, 28 ledger incidents



Each incident counted once by its first confirmed access path; attacker claims not counted. Classified from the ledger on pages 3 and 4.

What the data was

11

incidents exposed government identifiers: SSNs, national ID numbers, tax IDs, passport or license numbers

8

exposed health or patient data

6

exposed secrets rather than personal data: env vars, API keys, cloud credentials, tokens

Several incidents fall in more than one group.

The undisclosed category is a finding

Seven organizations, including Quest Diagnostics, IDScan.net and Aesto Health, had not said how the attacker got in by the end of October. For a security leader reading the notice, "an unauthorized party accessed certain records" is the whole root-cause analysis.

AI appeared where it never had

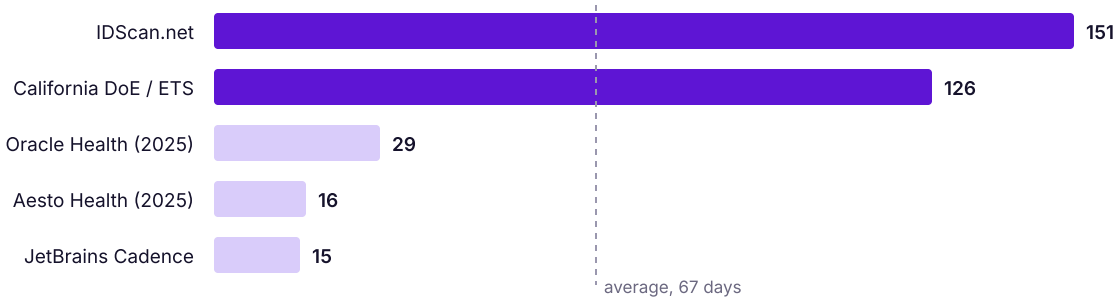
Attacker-side AI was named in four of the 28 incidents and in Microsoft's attribution of the ConnectWise intrusions. Employee-side AI misuse caused one more. A year earlier both counts were zero.

Dwell time, where disclosed

The five incidents with a disclosed access window and discovery date (IDScan, ETS, Oracle Health, Aesto Health, JetBrains Cadence) averaged about 67 days of attacker access before discovery. IDScan and ETS were each above four months.

> *Disclosure to affected people lagged discovery by anywhere from two weeks to six months.*

Days of attacker access before discovery, where both dates were disclosed



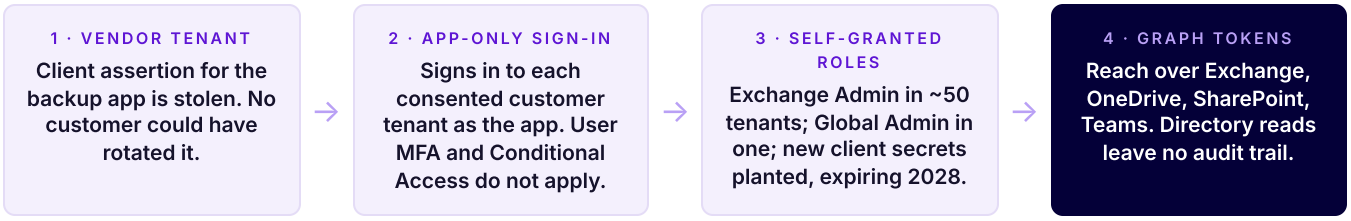
Access windows as disclosed by each organization; the ETS window ends at discovery, not at the last confirmed copy.

04 · TREND 1

The vendor is the *blast radius.*

Ten of the 28 incidents reached the victim through a third party's platform or credential. In every one, the third party held standing, consented access to sensitive data, and nothing was watching how that access was used.

How one stolen app credential reached ~1,000 Microsoft 365 tenants (ConnectWise Cloud Backup, Oct 15 to 17)



Huntress confirmed unauthorized sign-ins in 162 tenants by Oct 16 and sign-in activity across almost 1,000 Entra tenants, 300+ successful, by Oct 17. Microsoft attributed the activity to Storm-0940 and said the group used the open-source ARTEX AI framework to attempt sign-ins to more than 2,250 tenants.

The same structure, sector by sector

Veradigm said credentials stolen from an unnamed vendor's network were used against a customer-service API. The access was confined to that API, yet it was enough to pull patient PII including SSNs.

AdaptHealth attributed its 4.1 million-person breach to "a successful social engineering ploy that compromised the privileged account of a third-party contractor," after which the attacker walked through cloud patient management, document storage and EHR portals.

iRhythm's 8.2 million-person breach came through "third-party hosted business applications" reached by social engineering.

EY's exposure of Goldman Sachs and Man Group client tax data happened on a third-party IT service management platform where support tickets carried client tax files as attachments.

California's 5.7 million student records left through a contractor's SFTP site with a compromised login, copied over nearly five months.

Oracle Health's 19.9 million-person breach of legacy Cerner servers began, by the company's account, with compromised customer credentials.

Denmark's CPR register was drained not by breaking into the government but by abusing a private company's lawful lookup access. The anomaly that surfaced it was an unusually large bill for about 14 million searches.

Three things follow

The data was not stolen from a database the victim controlled; it was stolen from a workflow the victim had delegated. Support tickets, file-transfer drops, lookup APIs and backup apps are where sensitive data pools, and they rarely appear in a data inventory.

Identity controls built for humans did not fire. App-only tokens, vendor service accounts and federated identities (5,000 Dropbox accounts opened through a Lenovo ID flaw, no password needed) all bypass the MFA prompt.

The victims learned of the breach from someone else: a vendor's notice, a leak-site listing, an attacker's email.

> *The missing control is content-aware inspection of what a connected app or vendor identity actually reads and moves.*

05 · TREND 2

Secrets are *data.*

October produced three back-to-back incidents at developer platforms in which the prize was not customer PII but the credentials customers had stored to run their software. Each one turned a single compromise into thousands of downstream compromises.

Railway · Oct 20 to 21

Intruder decrypted stored environment variables for up to 1.85% of workspaces, about 100,000. Railway said the attacker was after AI provider API keys for model distillation, part of a campaign against platforms serving inference providers. Wallets with keys in env vars were drained two hours before disclosure.

Zeabur · Oct 14 to 15

"Long-lived administrative credentials that were stored in plaintext" on internet-reachable management interfaces. 3,218,965 env-var records with values, 566 plaintext API keys, root sign-ins on 101 servers. Admin tokens with no expiry dating to January 2024 had survived the August cleanup. Shutting down Dec 31.

JetBrains Cadence · Aug 8 to 24

TeamCity deserialization flaw (CVE-2026-63077), on CISA's KEV list since Aug 5, left unpatched. Attackers took a full 2024 server backup, AWS IAM users and secrets including employees', S3 files and possibly synced source. JetBrains: treat "any credentials or secrets stored in Cadence" as compromised.

Binance, about $6.34 billion: the chain by Binance's own account

JUNE	SEPTEMBER	OCT 20	OCT 29
Engineer A installs a tampered Homebrew package while searching for an AI coding tool	Engineer B runs a fake recruiter's "assignment" repo on a work machine; undetected	From B's machine, an internal AI tool is used to take over A's and C's; a container deployment token is found	Token plus stolen credentials deploy a container that reaches the wallet signing service

FBI attribution: North Korea's TraderTraitor cluster. About $2.85B frozen or intercepted. Binance has not disclosed what the internal AI tool was or how it was abused.

Two older patterns, same lesson

The help-desk vishing campaign Arctic Wolf tracks as **PREY-0058** phoned directors and vice presidents posing as IT, walked them through a look-alike Microsoft 365 login that harvested credentials and MFA approvals, then replayed the session tokens from residential proxies matching the victim's location. No malware, no lateral movement: SharePoint searches, then bulk pulls from SharePoint, OneDrive, Exchange and Box, then an extortion demand.

The **Denmark CPR** breach reportedly turned on a former employee's password of "123456," found in an unrelated data leak.

Why it keeps happening

A deployment token on a laptop, an AWS key in a CI backup, an API key in a .env file, a password in an old email: each is a sensitive record that was never classified as one, so nobody knew it existed until an attacker used it.

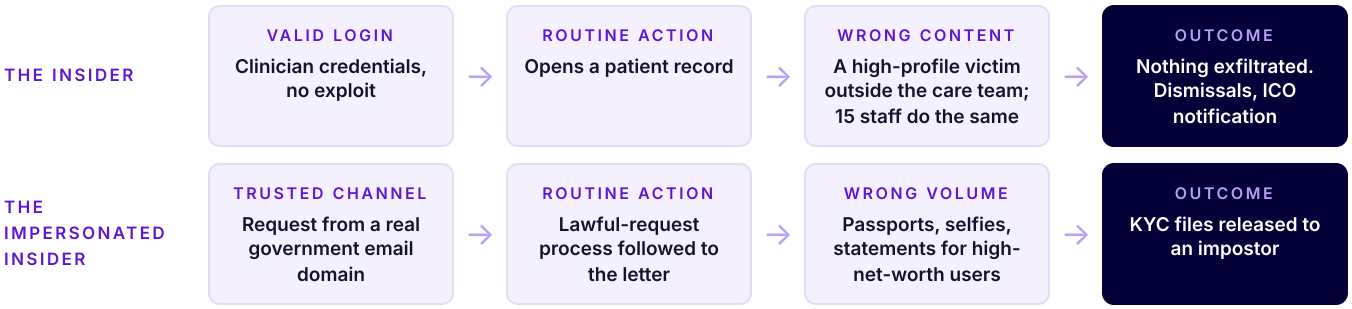
The practical response is twofold: discover secrets wherever they sit (code, chat, tickets, cloud storage, developer platforms) and stop them from being pasted, committed or uploaded in the first place. Rotation after the fact is the expensive path. The cheap path is never letting the key leave the vault.

06 · TREND 3

Valid identity, *wrong records.*

A cluster of incidents in which no system was exploited at all. Someone with legitimate access either misused it or handed it over, and the organization's controls could not tell the difference between that and ordinary work.

Two lanes, one failure: the action looked routine, the content and volume did not



The insider

At **Nottingham University Hospitals NHS Trust**, 15 staff, including doctors and nurses, were dismissed for opening the medical records of the June 2023 Nottingham attack victims with no clinical reason. Eleven were dismissed in May and the remainder later, with the BBC reporting the full tally in October.

At **East Suffolk and North Essex NHS Foundation Trust**, ten staff were suspended in late September over alleged access to a single patient's records.

The UK Health Secretary called the cases a disgrace. Both trusts have told regulators; Nottingham is testing a system to alert it to possible breaches. The records were never exfiltrated. They were read, by people whose job titles entitled them to read records in general, and that was enough to end careers and trigger ICO notifications.

The impersonated insider

Revolut disclosed on Sept 12 that it had handed passport and license copies, selfies, account statements and transaction histories to a party sending requests from a legitimate government agency's email domain. No breach of Revolut's systems occurred; a process for lawful data requests was followed to the letter against a fraudulent requester. ZachXBT reported the targets appeared to be high-net-worth users.

ASOS said an "unauthorized party" posed as a trusted contact to get an employee's work login, then used it on third-party platforms the retailer relies on. The attackers' most effective move was sending extortion threats as push notifications through the ASOS shopping app, which hit the share price before the company could respond.

Nikkei had a Google Workspace account accessed from late July and a Microsoft 365 account used Sept 30 to send about 9,000 phishing emails to staff and news sources, which in turn compromised a Nikkei BP employee.

The **Arizona court system** lost PII on about 1.3 million people after one employee clicked a phishing link on Sept 24. **DriveWealth** attributed a Sept 4 to 5 intrusion exposing customer names, employment details and portfolio values to "a social engineering campaign."

Content-aware monitoring of what a user actually touches is the control that separates misuse from work. The NHS trusts needed to know that specific patient records were being read by people outside the care team. Revolut needed a check on the sensitivity and scope of what was about to leave in a single response.

07 · TREND 4

AI arrived on *both sides.*

September and October were the months when AI stopped being a forecast in breach reporting and became a line item in the root cause. It showed up in three roles: the attacker's workforce, the victim's own tool misused, and the thing the attacker was trying to steal.

AI as the attacker's workforce

GreyNoise: hundreds of agents on OpenAI's Codex harness with a DeepSeek model built and ran exploits for two PaperCut flaws across 440 instances at 395 orgs in 48 countries. South Korean police arrested a former pentester who used the ARTEX framework against 90+ orgs for 8.2M records. Spain's AEPD logged its first breach notification attributing an incident to an autonomous agent.

AI as the victim's own tool

Singapore's PDPC: a Bee Cheng Hiang employee used a personal generative AI tool to write a bulk-email script; a missing bracket exposed 95,364 addresses. The regulator was explicit that the tool did not malfunction. Nobody reviewed the code, testing checked logs not the test email, and the company had no AI policy at all. Binance's pivot ran through "an internal AI tool."

AI as the target

Railway's intruder, by the company's account, was hunting AI provider API keys to run distillation attacks, and the campaign targeted platforms that serve inference providers. Symantec separately described an LLMjacking campaign run by an offensive AI agent. Stolen model access is now a commodity with its own supply chain.

How fast an agent swarm moves (PaperCut campaign, GreyNoise)

UNDER 4 HOURS
Empty workspace to first remote code execution on a real victim

→

+2 HOURS
First domain admin

→

26 SECONDS
11 organizations compromised once the campaign launched

→

7 MINUTES
A US high school, initial access to full domain admin

The operator told the agents to avoid eight countries. They did not consistently comply.

43%

of breaches studied involved shadow AI, up from 20% the year before; nearly half led to data loss

92%

of AI-related breaches involved no proper AI access controls; only 40% of organizations report having them

68%

of breached organizations had no AI governance, up from 63%; about a third of those with a policy had strict deployment approval

Source: IBM Cost of a Data Breach Report 2026.

The defensive implication is uncomfortable. Attacker agents compress the window between first access and bulk theft to minutes, and the victim's own AI tools generate code and move data that nobody reviews. Governance documents do not close that gap. What does is enforcement at the data layer: knowing what sensitive content is flowing into prompts, coding assistants and agent tool calls, and being able to block or redact it in line, synchronously, before the agent proceeds.

08 · TREND 5

Old data, *unwatched.*

The largest exposures of the period were PHI and PII that had sat in cloud archives, portals and transfer systems for years. In most cases the intrusion ran for months before anyone looked.

ORGANIZATION	DATA	PEOPLE	ACCESS WINDOW	DISCOVERED	DISCLOSED OR COUNTED
Oracle Health (legacy Cerner)	Names, SSNs, diagnoses, medications, lab results	19,929,149	Jan 22 to Apr 1, 2025	Feb 20, 2025	Count via Texas AG, Oct 2, 2026
Aesto Health	Names, SSNs, license numbers, financial accounts, medical and insurance data	9,540,683	Dec 2 to 18, 2025	Dec 18, 2025	Notice June 2026; HHS portal late Aug
iRhythm	Names, contact, insurance numbers, device serials, dates of service	8,179,772	Through June 8, 2026	June 8, 2026	Letters Oct 9, 2026
California DoE / ETS	Student names, DOBs, addresses, disability status, assessment data	5.7M+	May 22 to Sept 18, 2026	Sept 25, 2026	Oct 16, 2026
Florida Dept. of Education	SSNs, names, DOBs, addresses	~4.9M	Unknown	Oct 19, 2026	Oct 29, 2026
AdaptHealth	Names, contact, demographics, insurance, health information	4,115,802	June 5, 2026	July 2026	Count Sept 9, 2026
Arizona courts	PII, foster-care review reports, protective orders	~1.3M	Sept 24, 2026	Sept 24, 2026	Oct 7, 2026
IDScan.net	Names, DOBs, license, passport and ID numbers	Not disclosed	Apr 4 to Sept 2, 2026	Early Sept 2026	Sept 4, 2026
Quest Diagnostics	Patient records in the Quanum lab-ordering portal	Not disclosed	Unknown	Oct 21, 2026	8-K Oct 22, 2026

From first access to the public count: two timelines



Bar: access window. Open dot: discovery. Filled dots: public notice and public count. Oracle Health's count surfaced through the Texas attorney general's portal more than eighteen months after the breach.

Three things stand out

The data was old. Legacy Cerner servers, an EHR archive, student records back to 2006, foster-care files back to 2010. **It sat in a portal, not a database.** Quanum LSM, ETS's SFTP site and VeriScan hold regulated data because moving it is their function, and they are rarely in scope for the controls on the system of record. **Discovery was slow or external.** IDScan's intruder had five months; ETS's nearly five.

The consequences are already arriving

At least nine class actions were filed against IDScan within days. Oracle, Aesto, iRhythm and AdaptHealth are offering credit monitoring to millions. EY faces demands from Goldman Sachs for "objective evidence" that its fixes work. Veradigm's 8-K notes the company "has not determined the full extent of potential liabilities."

09 · WHAT WOULD HAVE CHANGED THE OUTCOME

Six controls, *none exotic.*

Each pattern points to a control that was absent at the moment the data moved. What they share is that they operate on the data itself, not on the network or the login.

DETECT	IDENTIFY	DECIDE	INVESTIGATE
What data is in this flow?	Who or what is moving it?	**Should it be allowed?**	What happened, what next?

Four questions, answered at the moment of movement. Each control below is one of them applied to a breach from this report.

1

Inventory the non-human identities that can read your sensitive data

The ConnectWise credential worked in every tenant where the app had been consented, and the app granted itself admin roles before anyone noticed. Enterprise app registrations, vendor service accounts, CI tokens and MCP connectors should be listed with the data each can reach, and client secrets added outside change control should alert.

2

Classify secrets as sensitive data and stop them at the point of exposure

Railway, Zeabur and JetBrains show that a platform breach becomes a thousand downstream breaches only because customers stored live keys there. Secrets pasted into chat, committed to repos, attached to tickets or saved to shared drives should be detected and blocked the same way an SSN would be.

3

Watch what delegated workflows actually carry

EY's exposure was a support ticket with a tax file attached. ETS's was an SFTP drop. Veradigm's was a customer-service API. Denmark's was a lookup service whose 14 million queries were only noticed on the invoice. Any system whose job is to move data for a third party needs content inspection and volume baselines.

4

Treat reads of regulated data as events, not logs

The NHS cases involved no exfiltration; the breach was the read. Clinical, HR, legal and finance teams need alerting when a record outside a person's working set is opened, and when many people open the same high-profile record. Compliance teams answering external requests need a gate on the sensitivity and volume of what is about to go out.

5

Put a control between AI tools and sensitive data, in both directions

Employees are writing code and moving data with AI tools whether or not a policy exists. Prompts, uploads, generated code and agent tool calls should pass through the same classification and enforcement as email and file sharing, with redaction or blocking in line. Because attacker agents compress theft to minutes, enforcement has to be synchronous, not a queue an analyst reviews tomorrow.

6

Shorten the gap between discovery and the people affected

Six months at Aesto Health and at EY, and more than eighteen before Oracle Health's count reached the public, are now the comparison points regulators and plaintiffs will use. Organizations that can say within days exactly which records left, and whose, are in a different legal position from those that cannot.

The outcome is a company that can keep adding vendors, agents and AI tools because it can prove, at any moment, *what touched its data and what was stopped.*

See it. Understand it. Stop it before it leaves.
Nightfall secures the first data flows in minutes, not quarters. Start in observe mode and turn on enforcement surface by surface.

nightfall.ai

10 · WHERE NIGHTFALL FITS

AI moves your data. *Nightfall controls it.*

Nightfall is the AI-native data security platform that gives enterprises real-time control over sensitive data movement across humans and AI agents, MCP servers, SaaS, email and endpoints. The breaches in this report are failures of exactly that control: the data was known to be sensitive, the path it left by was legitimate, and nothing in the path could tell work from theft.



PATTERN FROM SEPT TO OCT 2026	WHAT FAILED	NIGHTFALL CAPABILITY
Vendor apps and third-party credentials with standing SaaS access ConnectWise, Veradigm, AdaptHealth, iRhythm, EY, ETS	No content-aware inspection of what a consented app or vendor identity read and moved	Human exfiltration control across SaaS, email, endpoints and browsers. AI-native detection classifies the data in motion; enforcement blocks, redacts or coaches at the moment it is about to leave.
Secrets in developer platforms, backups, laptops and tickets Railway, Zeabur, JetBrains Cadence, Binance	Credentials were never treated as sensitive data, so no one knew where they sat until they were used	API keys, cloud credentials and tokens as first-class sensitive data types, blocked in chat, code, tickets, cloud storage and coding assistants before a secret is committed or shared.
Insider reads and impersonated requests NHS trusts, Revolut, ASOS, Nikkei	Valid identities doing routine-looking actions on the wrong records, at the wrong volume	Insider risk agents that triage, investigate and challenge each event in context, converting raw access into CISO-ready findings and recommended actions.
AI on both sides PaperCut, ARTEX, AEPD, Bee Cheng Hiang	No control between AI tools and sensitive data, no visibility into shadow AI	AI agent control that governs how copilots, coding tools and MCP-connected workflows access, process and move data. Enforcement is synchronous: the agent cannot proceed until the check returns. Shadow AI and agent chains are discovered.
Regulated data in archives, portals and transfer systems Oracle Health, Aesto Health, Quest, IDScan	Discovery depended on an outsider; the organization could not say what left or whose it was	Detection models trained on PHI, PII and financial data across SaaS and cloud storage, data lineage from origin to destination, and autonomous investigation that says exactly which records moved and when.

> *The unit of risk was a workflow, not a user or a file. Nightfall follows the data through the chain.*

SOURCES

Where the *facts come from.*

Pages opened for this report, by incident. Roundups (Strobes, SOCRadar, CM-Alliance) were used to locate incidents and, where noted in the ledger, as the source of record for incidents whose primary notice could not be opened.

Third-party and SaaS access. Huntress, "ConnectWise (SkyKick) vendor credentials abused in M365 attacks," Oct 16 to 21, 2026. The Register, "Microsoft says ConnectWise breach let Chinese spies loose on 2,250 cloud tenants," Oct 22, 2026. NetSec.News, "Veradigm reports third-party data breach," Sept 15, 2026. BleepingComputer, "AdaptHealth confirms 4.1 million people exposed," Sept 9, 2026. HIPAA Journal, "iRhythm data breach affects 8.2 million," Oct 19, 2026. Cyber Security News, "EY data breach exposes Goldman Sachs and Man Group clients' data," Oct 7, 2026. Cybernews, "Millions of California student records breached," Oct 2026. Malwarebytes, "How '123456' became the key to almost 9 million personal records," Oct 12, 2026. BleepingComputer, "Dropbox accounts breached through Lenovo email verification flaw," Sept 2, 2026.

Secrets and credentials. The Hacker News, "Railway says intruder may have accessed customer secrets," Oct 2026. Cybernews, "Zeabur shutting down after two back-to-back breaches," Oct 2026. The Hacker News, "Attackers breached JetBrains Cadence via unpatched TeamCity," Sept 5, 2026. StepSecurity, "Binance's $6.34 billion hack: developer machines, AI tooling, and wallet signing," Nov 2026. The Hacker News, "Fake IT calls target executives in Microsoft 365 data theft and extortion attacks" (Arctic Wolf, PREY-0058), Sept 7, 2026.

Insiders and social engineering. Paubox, "NHS trust dismisses 15 staff for viewing attack victims' records," Nov 10, 2026. TechCrunch, "Revolut confirms customer data breach through fake government requests," Sept 12, 2026. Bloomberg via Insurance Journal, "Hackers tricked ASOS employee to steal work account login," Oct 8, 2026. Infosecurity Magazine, "Nikkei discloses two employee cloud account compromises," Oct 6, 2026. Bloomberg Law, "Arizona court cyberattack exposes data on 1.3 million people," Oct 7, 2026. Law360, "Squire Patton Boggs, Holland & Knight report data breaches," Oct 9, 2026.

AI on both sides. BleepingComputer, "AI-powered attack exploited PaperCut flaws to hack 395 organizations" (GreyNoise), Sept 10, 2026. BleepingComputer, "Former pentester used ARTEX AI to steal 8.2 million data records," Oct 26, 2026. Help Net Security, "Spain reports first data breach involving autonomous AI agent," Sept 17, 2026. Cyber Security News, "Chainalysis confirms cyberattack using agentic techniques," Oct 31, 2026, and the company's Nov 4 update as reported. Frontier Enterprise, "Shadow AI risks: Bee Cheng Hiang's email leak" (PDPC), Oct 21, 2026. IBM, Cost of a Data Breach Report 2026, figures as summarized by Varonis.

Regulated data. SecurityWeek, "9.5 million impacted by Aesto Health data breach," Sept 1, 2026. Rescana, "Quest Diagnostics Quanum LSM incident," Oct 25, 2026. Strobes, "October 2026 data breaches and key lessons" (Oracle Health count via Texas AG). Strobes, "Biggest data breaches of September 2026" (IDScan.net notice, DriveWealth, BigCommerce/Ribon).

Roundups used for discovery. SOCRadar, major cyber attacks of September and October 2026. CM-Alliance, data breach tables for September and October 2026.

Nightfall AI. Product and platform descriptions reflect Nightfall's current platform and messaging as of November 2026.

© 2026 Nightfall AI. All rights reserved. Third-party names and marks are the property of their respective owners. Incident facts are reported as published by the cited sources as of November 12, 2026 and have not been independently verified by Nightfall. Attacker claims are labeled as such and should not be read as confirmed figures.