

SIGNALS 2026

Five signals that cut through the noise.

Every vendor report says its problem is the biggest. We read two years of them, checked every number against the source, and kept only what independent sources agree on, what has lasted, and what changes a board decision.

FIR Risk Advisory
For CEOs and boards · November 2026
firrisk.ai

Most of what you read about cyber risk this year was noise. This report is what's left.

Every year the security industry publishes hundreds of reports. Most are careful work, and each one makes the case that its problem is the biggest. For a CEO or a board, the result is noise.

Over the past two years, FIR Risk has published 100 editions analyzing the major reports, from Verizon, Microsoft, Mandiant, CrowdStrike, IBM, Anthropic and many others. This year we asked a different question: **what survives when you put them all side by side?**

SIGNALS is the answer: five findings that independent sources agree on, that have held for more than a year, and that matter to the decisions leaders make about priorities, funding and oversight. We went back to the original reports and checked every number. What didn't pass, we left out.

Each signal ends with one question for your next board meeting. They are not mandates. They are the questions we would want answered if it were our company.

Page 3 puts all five signals on one page. Pages 5 to 9 take them one at a time. Page 14 gives you the five questions to take into your boardroom. This year you can also listen: a narrated edition of SIGNALS is at firrisk.ai.

Thank you for reading FIR Risk this year.

Bruce Bird
FIR Risk Advisory

THE WHOLE PICTURE

The five signals on one page

If you read nothing else: what held up across two years of threat reports, the number that proves it, and the question it puts to the board.

01 The window is gone. Expose less.

31%

of breaches start with an exploited vulnerability, now the #1 way in. Up from 20%.

ASK

Do we know every system we expose to the internet and every open-source package we depend on, and is that footprint getting smaller?

02 Inside, they look like you.

82%

of intrusions CrowdStrike detected in 2025 used no malware. 71% in 2022.

ASK

If an attacker were logged in as one of our employees today, would we know?

03 Half of breaches now involve a third party.

48%

of breaches involved a third party. 9% three editions ago.

ASK

How many vendors do we rely on, and how many of the critical ones have a current risk assessment?

04 Recovery time sets the bill.

~10×

the economic loss when restoring takes almost two weeks instead of under one.

ASK

How quickly could we really restore core operations, and have we tested it?

05 Your own AI is the risk you control.

45%

of employees now use AI regularly at work, up from 15% a year earlier.

ASK

Which AI tools and agents are in use across our company, approved or not, and who owns each one?

Sources: Verizon DBIR 2023–2026 · CrowdStrike Global Threat Report 2024–2026 · Microsoft/Accenture, Securing Nations in the Intelligent Economy (2026). Every figure is checked against the primary report; the full list, with page references, is in the web edition at firrisk.ai.

METHOD

How we filtered

We started with everything we published and kept only what passed four tests.

100 FIR Risk editions, and counting

26 primary reports re-read · 19 made the cut

43 key figures checked at the source page

5 signals

TEST 1

Agreement

At least three independent sources, using different data, reach the same finding.

TEST 2

Persistence

It held across more than one year of data, not a single season's headline.

TEST 3

Direction

The trend is clear: we can say which way it is moving, and how fast.

TEST 4

Decision value

It matters to decisions a CEO or board makes: priorities, funding or oversight.

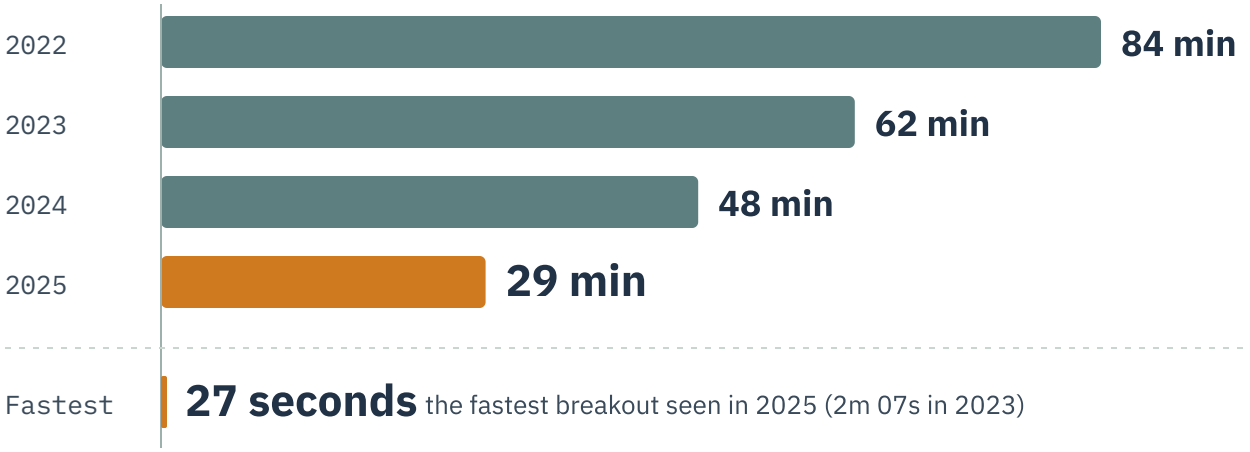
Every number in this report traces to a page in a primary report and carries the year of its data.

What failed the tests is on page 11, Noise we filtered. Long-fuse risks like quantum can't pass a persistence test until it's too late, so we track them on page 13, Weak signals to watch.

The window is gone. Expose less.

Attackers now move faster than any patch cycle. The lever left to you is how much you expose.

Time from first access to moving through the network, average



Average eCrime "breakout time." CrowdStrike Global Threat Report, 2024–2026 editions.

31%

of breaches now start with an exploited vulnerability, the #1 way in.

Verizon DBIR 2026

29%

of exploited flaws were attacked on or before the day they were published.

VulnCheck 2026

26%

of critical known-exploited flaws were fully fixed in 2025, down from 38%.

Verizon DBIR 2026

ALSO AGREE Mandiant · exploits #1 way in, 32% · IBM X-Force · app exploits up 44% · Microsoft · weaponized in under 24 hours

You can't out-patch the window. Shrink what you expose, starting with internet-facing systems and open-source packages.

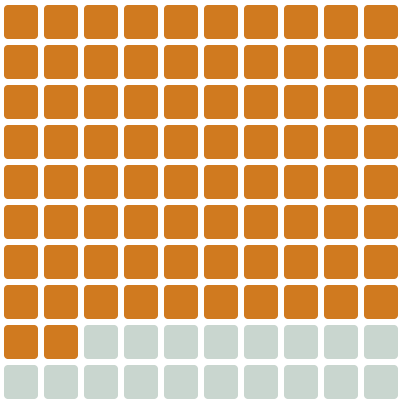
ASK Do we know every system we expose to the internet and every open-source package we depend on, and is that footprint getting smaller?

FIR reading: INTEL-30 The Pre-Disclosure Window · INTEL-32 When the Sellers Agree · INTEL-45 Stop Counting Patches · E91 · E94

Inside, they look like you.

Logins are harder to break, so attackers talk their way in, then act like real users.

Of every 100 intrusions CrowdStrike detected in 2025

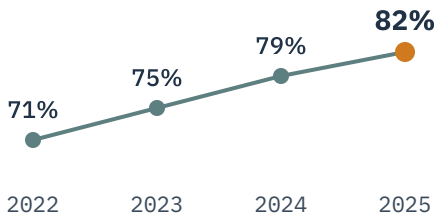


82

used **no malware**: stolen logins, real tools, normal-looking activity

18 used malware your tools are built to catch

RISING EVERY YEAR



Malware-free share of detections. CrowdStrike Global Threat Report, 2024–2026 editions.

#2

Voice phishing is now the second most common way in. Email phishing fell from 14% to 6%.

Mandiant M-Trends 2026

+502%

malicious voice calls over Microsoft Teams in a year.

Microsoft Digital Defense Report 2026

83%

of privilege escalation used no exploit, only permissions that already existed.

Verizon DBIR 2026

ALSO AGREE · Unit 42 · identity a factor in ~90% of cases · Microsoft · password attacks down 26%

A valid login is the best disguise. Know who has access, and notice when they act out of character.

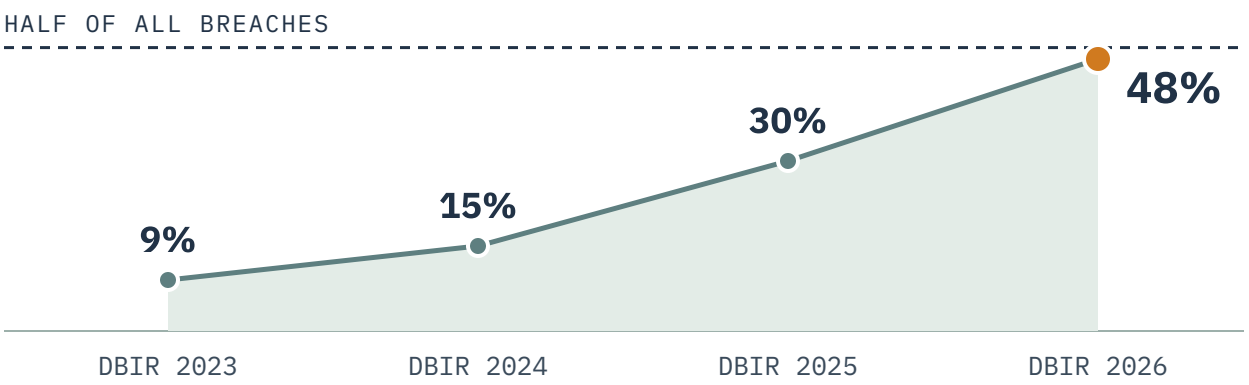
ASK If an attacker were logged in as one of our employees today, would we know?

FIR reading: INTEL-8 The Breach Starts With a Phone Call · INTEL-23 The Permission Problem · INTEL-42 The First Email Asks for Nothing · E85 · E94

Half of breaches now involve a third party.

48% of breaches involved a vendor, partner or supplier, up from 9% three editions ago.

Breaches involving a third party, now close to half



Verizon Data Breach Investigations Report, 2023–2026 editions (each covers the prior year).

23%

of third parties fully fixed missing MFA on their cloud accounts.

Verizon DBIR 2026

~4×

more major supply-chain and third-party breaches over five years.

IBM X-Force 2026

~200

customers' data taken after one SaaS provider was breached.

Anthropic, Sept 2026

ALSO AGREE

- Anthropic · one stolen token → cloud admin in ~3 hours
- Cyentia · trusted relationships: 20% of extreme events, 65% of losses
- Mandiant · unmonitored OAuth grants and broad API permissions

That 48% is an average. Your number depends on which vendors hold your keys, so find out.

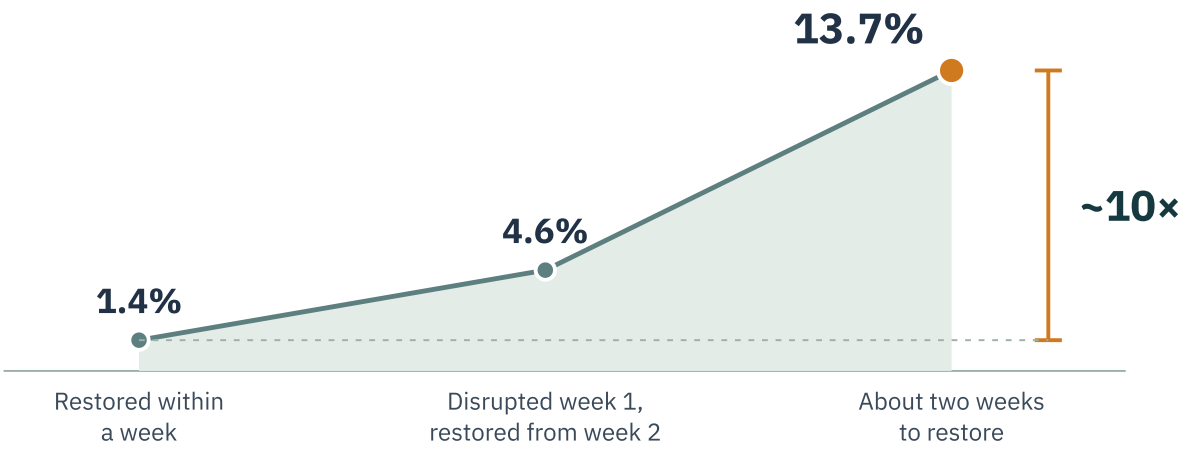
ASK How many vendors do we rely on, and how many of the critical ones have a current risk assessment?

FIR reading: INTEL-25 The Third-Party Half · INTEL-31 Access Is a Product · INTEL-38 Three Hours to Admin · E90 · E93

Recovery time sets the bill.

Fewer victims pay, so attackers go after your ability to recover. The cost is now set by how long you're down.

Economic loss grows much faster than downtime



Modeled loss, % of monthly GDP, for a cyberattack on one critical sector (upstream oil and gas in an oil-dependent economy). Microsoft / Accenture, Securing Nations in the Intelligent Economy, 2026, p. 20.

69%

of ransomware victims refused to pay, up from about half in 2022.

Verizon DBIR 2025 and 2026

£1.9B

Jaguar Land Rover: a five-week stoppage, the UK's costliest cyberattack.

Verizon DBIR 2026

+63%

data-extortion events in 2025. Fewer payments haven't meant fewer attacks.

Intel 471, 2026 Outlook

ALSO AGREE · Mandiant · backups, identity and hypervisors hit first · DBIR · ransomware in 48% of breaches · Cyentia · disruption in 43% of extreme events, 8% of all

Paying is no longer the big decision. Recovery speed is.

ASK How quickly could we really restore core operations, and have we tested it?

FIR reading: INTEL-17 The Ecosystem Is Eating Itself · INTEL-28 The Ransom Bluff · INTEL-33 The Cost of Week Two · E85 · E90 · E92

Your own AI is the risk you control.

Attackers' AI makes old attacks cheaper. Your own AI use tripled in a year, much of it on personal accounts. That part you control.

Two numbers on the same 0–100% scale

WHAT GETS THE ATTENTION: ATTACKERS' AI

of 793 threat actors caught using AI were rated high or critical risk · Verizon DBIR 2026, with Anthropic data

<1%

WHAT YOU CONTROL: YOUR OWN AI

of employees use AI regularly on work devices; 67% of AI users sign in with personal accounts · Verizon DBIR 2026



45%

What AI really changed: cost and speed. In one espionage campaign, AI did 80–90% of the work. Anthropic, Nov 2025

What it hasn't changed: phishing success. Attacks got cheaper and faster, not more likely to land. Verizon DBIR 2026

20,000+

enterprise tenants reached by fake AI browser extensions that harvested chat histories.

Microsoft Digital Defense Report 2026

+$670K

added to a breach's cost when unsanctioned AI was involved.

IBM Cost of a Data Breach 2025

None

of the three leading AI labs expects prompt injection to be fully solved.

OpenAI, Anthropic, Google, 2026

ALSO AGREE SEC filings · Microsoft, Alphabet, Apple now list their own AI agents as a risk (Q3 2026)

Watch the attackers' AI. Govern your own.

ASK Which AI tools and agents are in use across our company, approved or not, and who owns each one?

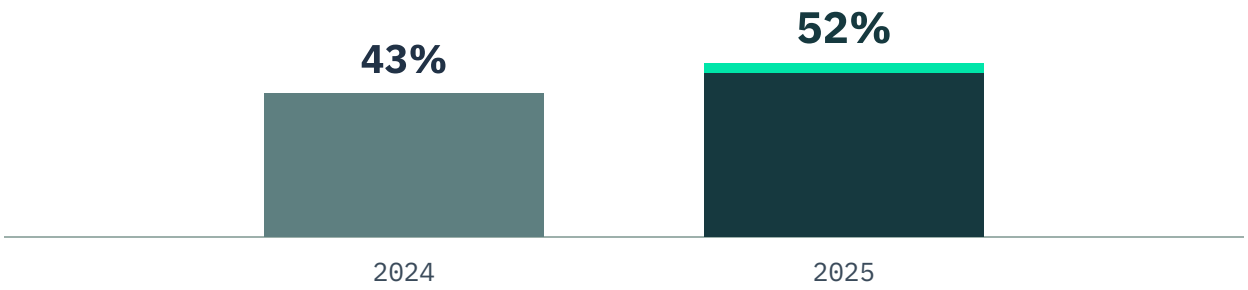
A good answer names an owner who can say what each agent can read, and what it is allowed to do.

FIR reading: INTEL-24 Breadth, Not Depth · INTEL-26 The Convenient Insider · INTEL-41 · INTEL-46 Know Your Agents · E93 · E95

The basics are working.

The controls everyone already knows still stop most attacks, and defenders are catching more on their own.

Breaches organizations found themselves, rather than being told



Share of compromises detected internally. Mandiant M-Trends 2026.

>99%

of unauthorized sign-in attempts are blocked by MFA.

Microsoft Digital Defense Report 2025

−26%

password attacks in a year, as MFA and passkeys force attackers elsewhere.

Microsoft Digital Defense Report 2026

>90%

of breaches were enabled by gaps that were preventable. The fixes are known.

Unit 42 Incident Response Report 2026

ALSO AGREE DBIR · 69% of ransomware victims refused to pay

The fundamentals still decide most outcomes. Keep funding them, consistently.

BEFORE ALL FIVE Which basics are we not funding yet?

FIR reading: INTEL-12 The Defender's Window · INTEL-36 The Alarm Is Yours · INTEL-44 The Basics Are Working · E85 · E94

FILTERED OUT

The noise, and what the data says

These claims were loud this year. Each one failed at least one of our tests: independent sources didn't agree, it didn't last, or it doesn't matter to the decisions leaders make.

WHAT YOU HEARD	WHAT THE EVIDENCE SHOWS
"AI super-attacks are here."	Fewer than 1% of 793 threat actors caught using AI were rated high or critical risk, and phishing's success hasn't moved. AI made attacks cheaper and faster; it hasn't made them land more often. Verizon DBIR 2026, with Anthropic data
"Just patch faster."	29% of vulnerabilities known to be exploited in 2025 were attacked on or before the day they were published. Speed can't win that race; reducing what you expose can. VulnCheck 2026 · Microsoft Digital Defense Report 2026
"The leak site lists who paid."	For a typical ransomware group, only about 9% of the victims it names publicly ever paid. The list is a pressure tactic, not a scoreboard. Verizon DBIR 2026
"Ransomware is fading."	Payments are fading: 69% of victims didn't pay. Attacks are not. Ransomware appeared in 48% of breaches, and data-extortion events rose about 63%. Verizon DBIR 2026 · Intel 471, 2026 Outlook
"Billions of attacks blocked."	Volume counts measure how much a vendor can see, not how much risk you carry. We left them out of this report. FIR view
"A breach costs $4.44M."	An average hides the tail. In the costliest 5% of events, the median loss was $100M, and the most common outcome was disruption, not stolen data. See Signal 04. Cyentia IRIS Xtreme 2026 · IBM Cost of a Data Breach 2025

Tested for agreement, persistence, direction and decision value.

How our past calls held up

We graded the threats we named in our 2024 and 2025 reports against this year's evidence.

4 held | **1** changed shape | **1** partly right

Business email fraud 2024 report	HELD	46 million+ business-impersonation attempts in 12 months, and the first message now rarely asks for money. Microsoft 2026
Ransomware and extortion 2025 report, #1	CHANGED SHAPE	Payments fell (69% refused) while extortion events rose 63%. Attackers now go after recovery. See Signal 04. DBIR 2026 · Intel 471
Infostealers and stolen credentials 2025 report, #2	HELD	Stolen identities remain central once attackers are inside: 82% of intrusions used no malware. See Signal 02. CrowdStrike 2026
Cloud compromise 2025 report, #3	HELD	The route now runs through vendors and tokens: one stolen token reached cloud admin in about three hours. See Signal 03. Anthropic 2026
AI-powered attacks 2025 report, #4	PARTLY RIGHT	AI made attacks cheaper and faster, but fewer than 1% of AI-using actors rated high-risk. See Signal 05. DBIR 2026
Insiders and fake IT workers 2025 report, #5	HELD	Fake remote IT workers drew on an estimated 15,000 stolen identities. DBIR 2026

Weak signals to watch in 2027

Real, but not yet proven across sources or years. Any of these could be next year's signal.

Attacks run by AI agents

In one espionage campaign, AI did 80–90% of the work, with people stepping in at a handful of decision points.

Not yet a signal: one documented case.

Anthropic, Nov 2025

AI on the defender's side

AI labs and security vendors have formed coalitions to find and fix vulnerabilities with frontier models.

Not yet a signal: too early to show up in breach data.

Anthropic Project Glasswing, Apr 2026

Post-quantum deadlines

US federal systems are to move to post-quantum encryption by 2035; EU and UK critical infrastructure by 2030.

Not yet a signal: the deadlines are set; the threat date isn't.

Microsoft / Accenture, 2026

Intruders staying unseen longer

Median time to detection rose from 11 to 14 days. Microsoft reports the same direction.

Not yet a signal: partly driven by long-running espionage cases.

Mandiant M-Trends 2026 · Microsoft 2026

Prompt injection

Practitioners rank it the top risk for AI applications, and none of the three leading AI labs expects it to be fully solved.

Not yet a signal: it barely appears in incident data so far.

OWASP · CrowdStrike 2026 · Microsoft Digital Defense Report 2026

Every signal starts as one report. We'll tell you when these cross the line.

Follow along weekly in FIR Risk Tuesday and FIR Risk INTEL at firrisk.ai.

Five questions for your next board meeting

Five signals held up across two years of independent evidence. Each one comes down to a single question.

01

The window is gone. Expose less.

Do we know every system we expose to the internet and every open-source package we depend on, and is that footprint getting smaller?

02

Inside, they look like you.

If an attacker were logged in as one of our employees today, would we know?

03

Half of breaches now involve a third party.

How many vendors do we rely on, and how many of the critical ones have a current risk assessment?

04

Recovery time sets the bill.

How quickly could we really restore core operations, and have we tested it?

05

Your own AI is the risk you control.

Which AI tools and agents are in use across our company, approved or not, and who owns each one?

And the basics still work. MFA blocks more than 99% of unauthorized sign-ins.

BEFORE ALL FIVE Which basics are we not funding yet?

MFA figure: Microsoft Digital Defense Report 2025.

APPENDIX

The evidence behind each signal

Which independent sources reached each finding, and how each signal scored on the four tests.

SIGNAL	Verizon DBIR	CrowdStrike	Mandiant	Microsoft	IBM	Unit 42	VulnCheck	Intel 471	Anthropic	OpenAI · Google	Sources
01 **The window is gone**	●	●	●	●	●	·	●	·	·	·	**6**
02 **Inside, they look like you**	●	●	●	●	·	●	·	·	·	·	**5**
03 **Third parties**	●	·	●	·	●	·	·	·	●	·	**5**
04 **Recovery time**	●	·	●	●	·	·	·	●	·	·	**5**
05 **AI**	●	·	·	●	●	·	·	·	●	●	**6**

Microsoft includes the Microsoft / Accenture economic model. IBM includes X-Force and Cost of a Data Breach. Signal 05 counts OpenAI and Google separately. Signals 03 and 04 also count Cyentia IRIS Xtreme 2026.

SIGNAL	PERSISTENCE	DIRECTION	DECISION IT INFORMS
01 The window is gone	2022–2025	Faster every year: breakout 84 → 29 minutes	Shrinking the internet-facing and open-source footprint
02 Inside, they look like you	2022–2025	Rising: malware-free 71% → 82%	Identity, access and user-behavior monitoring
03 Third parties	4 DBIR editions	Rising every year: 9% → 48%	Vendor risk assessment coverage
04 Recovery time	2022–2025	Non-payment ~50% → 69%; extortion still up	Funding tested recovery
05 AI	2025–2026	Rising: AI use at work 15% → 45%	Ownership of AI tools and agents

Every signal passed all four tests. One sits near the bar, and we say so: Signal 05 is the youngest, with two years of data. Signal 03 started at the minimum three sources; Mandiant and Cyentia have since reached the same finding. Sources are on page 18; every figure, with its source page, is in the web edition at firrisk.ai.

The five signals in MITRE ATT&CK

Hand these two pages to your security leader. Each signal is translated into ATT&CK, the framework most security teams work in: the techniques behind it, how often Mandiant saw them, the mitigations MITRE lists for them, and one question about coverage.

01

The window is gone

Initial Access

Attackers exploit internet-facing applications and remote-access services, often within a day of a flaw becoming known.

T1190	Exploit Public-Facing Application	27.7%
T1133	External Remote Services	23.8%

MITRE MITIGATIONS

- M1035 Limit Access to Resource Over Network
- M1042 Disable or Remove Feature or Program
- M1030 Network Segmentation
- M1016 Vulnerability Scanning

> **Which internet-facing systems would we know were exploited within hours, not days?**

02

Inside, they look like you

Initial Access · Stealth

A convincing call or a stolen login, then activity that looks like a normal user. ATT&CK v20 names this tactic Stealth: “appearing as normal behavior.”

T1078	Valid Accounts	17.0%
T1566.004	Spearphishing Voice	6.1%
T1684.001	Impersonation (T1656 in M-Trends)	1.7%

MITRE MITIGATIONS

- M1032 Multi-factor Authentication
- M1036 Account Use Policies
- M1026 Privileged Account Management
- M1017 User Training

> **How much valid-account activity can we tell apart from the real user?**

03

Third parties

Initial Access · Lateral Movement

A vendor’s access, software or token used to reach you. One SaaS breach reached about 200 of its customers.

T1550.001	Application Access Token	5.3%
T1195	Supply Chain Compromise	1.7%
T1199	Trusted Relationship	0.5%

Rare as techniques in Mandiant’s cases, yet present in 48% of DBIR breaches: Mandiant counts techniques, Verizon counts breaches.

MITRE MITIGATIONS

- M1018 User Account Management
- M1032 Multi-factor Authentication
- M1036 Account Use Policies
- M1021 Restrict Web-Based Content

> **Would we see a vendor’s token used from somewhere new, and could we cut it off?**

Bars: share of Mandiant’s 2025 investigations in which the technique was observed (M-Trends 2026, pp. 85–101). Mitigations are a selection from those MITRE lists for each technique.

The five signals in MITRE ATT&CK, continued

04

Recovery time

Impact

Backups and recovery tools are disabled first, then data is encrypted, stolen, or both, to force a payment.

T1486	Data Encrypted for Impact	17.7%
T1657	Financial Theft	10.4%
T1490	Inhibit System Recovery	2.7%

Encryption dropped out of Mandiant's top 10 as Data Staged (T1074, 39.6%) rose: a shift from locking data to stealing it.

MITRE MITIGATIONS

- M1053 Data Backup
- M1028 Operating System Configuration
- M1018 User Account Management
- M1040 Behavior Prevention on Endpoint

> **Would we detect backups or recovery tools being disabled before encryption starts?**

05a

AI, attackers' side

Resource Development

Attackers use AI to work faster on familiar techniques. Fewer than 2.5% of the AI-assisted techniques Verizon observed were rare ones; for initial access, phishing (44%) and exploitation (32%) led.

T1588.007	Obtain Capabilities: Artificial Intelligence	0.2%

MITRE: this "cannot be easily mitigated with preventive controls." The defense is the controls on Signals 01 to 03, which AI speeds up.

MITRE MITIGATIONS

- M1056 Pre-compromise

> **Which of our existing controls get weaker if attacks get cheaper and faster?**

05b

AI, your side

MITRE ATLAS · Execution

Your own AI tools and agents can be steered by text they read. MITRE ATLAS catalogs attacks on AI systems the way ATT&CK catalogs attacks on networks.

AML.T0051	LLM Prompt Injection

ATLAS MITIGATIONS

- AML.M0019 Control Access to AI Models and Data in Production
- AML.M0020 Generative AI Guardrails
- AML.M0033 Input and Output Validation for AI Agent Components
- AML.M0024 AI Telemetry Logging

> **Which of our AI tools and agents are in our threat model, and who owns each one?**

Techniques and mitigations: MITRE ATT&CK v20 (October 2026) and MITRE ATLAS 5.6. Frequencies: M-Trends 2026, pp. 40–41, 85–101. AI technique research: Verizon DBIR 2026, pp. 26–27, with Anthropic data. Mitigations are MITRE's listed options, shown to inform discussion with your security team, not as prescriptions.

Where every number comes from

100 FIR Risk editions → **26** primary reports re-read → **19** cited → **43** figures checked → **5** signals

SOURCE · REPORT	01	02	03	04	05	2027
Anthropic · AI-orchestrated espionage campaign (Nov 2025) · Detecting and countering misuse of AI (Sept 2026) · Project Glasswing (Apr 2026)	·	·	●	·	●	●
CrowdStrike · Global Threat Report, 2024–2026 editions	●	●	·	·	·	●
Cyentia Institute · IRIS Xtreme 2026	·	·	●	●	·	·
IBM · Cost of a Data Breach Report 2025	·	·	·	·	●	·
IBM X-Force · Threat Intelligence Index 2026	●	·	●	·	·	·
Intel 471 · Cyber Threat Trends & Outlook 2026	·	·	·	●	·	·
Mandiant (Google Cloud) · M-Trends 2025 and 2026	●	●	●	●	·	●
Microsoft · Digital Defense Report 2025 and 2026	●	●	·	·	●	●
Microsoft & Accenture · Securing Nations in the Intelligent Economy (2026)	·	·	·	●	·	●
Microsoft, Alphabet, Apple · Q3 2026 10-Q and 10-K filings (SEC)	·	·	·	·	●	·
OpenAI, Anthropic, Google · AI agent security guidance (2026)	·	·	·	·	●	●
OWASP · Top 10 for LLM Applications 2026 · Top 10 for Agentic Applications	·	·	·	·	·	●
Palo Alto Networks Unit 42 · Global Incident Response Report 2026	·	●	·	·	·	·
Verizon · Data Breach Investigations Report, 2024–2026 editions	●	●	●	●	●	·
VulnCheck · Exploit Intelligence Report 2026	●	·	·	·	·	·

01–05: the signal pages each source supports. 2027: the watch list on page 13. Unit 42, OWASP, Project Glasswing and the AI labs' guidance were checked against the publishers' web pages.

Also re-read. Red Canary Threat Detection Report 2026 · Rapid7 Global Threat Landscape Report 2026 · Picus Red Report 2026 · GuidePoint GRIT 2026 Ransomware and Cyber Threat Report · Zscaler ThreatLabz 2026 AI Security Report · Anthropic 2026 Agentic Coding Trends Report and 2026 State of AI Agents Report. Useful context; no figure in this report depends on them.

The FIR library. 100 FIR Risk editions, and counting. The web edition at firrisk.ai lists the editions behind each signal, and every figure with its source page.

How to read the numbers. Each figure is quoted as its source reports it, for the year its data covers. Sources use different datasets and definitions, so figures from different reports are not added or compared directly.

About this report. SIGNALS is general risk intelligence for leaders. It is not legal, compliance or security advice for a specific organization.

FIR Risk Advisory
Verified risk intelligence for CEOs and boards.
firrisk.ai