



HumanRiskShield

HUMAN RISK MANAGEMENT

INDEPENDENT SYNTHESIS | VERIFIED SOURCES | 2026

The Human Risk Report

The risk lives in the handoff.

People are not the weakest link. Unprotected decisions are. This report maps where people, identities, workflows, AI and partners meet - and the controls that keep one bad moment from becoming a material event.

VERIFIED SIGNAL

62%

of breaches had a human element present

Not the same as cause or initial access.

[V26] N=22,624 BREACHES



RESEARCH CUT-OFF 24 AUGUST 2026

SOURCES [\[V26\] VERIZON 2026 DBIR](#)

THE 2026 PICTURE

Six numbers. Six decisions.

The threat is not only the message someone clicks. It is the identity behind it, the workflow that trusts it, the supplier connected to it and the speed of the response.

62%

Human element present

Design safer decisions; do not blame people.

[V26]

31%

Vulnerabilities led entry

Patch exposed systems and protect identities together.

[V26]

48%

Third-party involvement

Treat supplier access as part of your own control plane.

[V26]

\$20.9B

Reported IC3 losses

Make fraud response contacts visible before an event.

[F25]

\$3.05B

BEC reported losses

Verify payment changes on a known second channel.

[F25]

58%

Attempted dollars frozen

Fast bank and IC3 reporting can change the outcome.

[F25]

CONTEXT, NOT FORECAST

\$4.99M

IBM's 2026 global average breach cost is a macro benchmark from 602 already-breached organizations. It is not an expected SMB loss. Use it to understand cost pressure, not to price your own risk.

HOW TO READ THE DATA

A number is only useful with its denominator.

This rebuild separates breach evidence, reported complaints, cost studies, vendor benchmarks and guidance. They answer different questions and must not be added together.


BREACH DATA

Observed or reported incidents and confirmed breaches. Strong for patterns; not a random probability sample.

[V26]


COMPLAINT DATA

Victim reports and reported losses. Underreported, sometimes overlapping and not always independently verified.

[F25]


COST STUDY

Interview-based averages from organizations that already had a breach. Not a forecast for one SMB.

[I26]


VENDOR BENCHMARK

Customer simulation telemetry. Useful for comparison; affected by selection, program maturity and scenario design.

[K26]

WORDS MATTER
PRESENT IN

A factor appeared somewhere in the event; it does not prove root cause.

INITIAL ACCESS

The first observed way in; later attack-chain steps may differ.

REPORTED LOSS

What complainants reported; not total economic harm.

SIMULATION

A controlled exercise; not a real-world compromise rate.

PUBLISHER INCONSISTENCY HANDLED

33.2 -> 20.1 -> 4.2


KnowBe4's 2026 materials describe the same trajectory as both a 79% and an 87% reduction. This report presents the raw percentages, dataset size and caveat instead of repeating the conflicting reduction claim.

THE 2026 SHIFT

Technology opened the path. The handoff decided the damage.

Vulnerability exploitation became the leading initial entry point. Human decisions, identity, suppliers and response still shaped how far the event traveled.

SELECTED INITIAL ACCESS VECTORS



Categories describe first observed access in non-error, non-misuse breaches. Method changes affect year-to-year comparisons.

62%

Human element present
Across the full breach dataset.

[V26]

48%

Third-party involvement
Up 60% year over year.

[V26]

THE PRACTICAL CONCLUSION

Training-only programs are too narrow.

Human risk lives across five connected pathways. Each one needs behavior change and a control that can stop a single mistake from becoming a material event.

01 IDENTITY

02 CHANNELS

03 PAYMENTS

04 AI + DATA

05 PARTNERS

PATHWAY 01

A valid login can make an attacker look ordinary.

Credential abuse fell as the first observed entry point, partly because Verizon changed its treatment of pretexting. Across the full attack chain, credentials remained a leading chokepoint.

39%

Credential abuse anywhere

Across the breach progression.

[V26]

13%

Credential abuse at entry

16% under the prior comparison method.

[V26]

97%

Password-spray share

Microsoft identity-attack telemetry; not breach share.

[MS25]

THE ATTACK CHAIN



CONTROLS THAT CHANGE THE PATH

- ✓ Use FIDO2/WebAuthn passkeys or security keys for admin, email, banking and remote access.
- ✓ Remove dormant accounts; review privileges and supplier access on a set cadence.
- ✓ Monitor exposed credentials and revoke active sessions, not just passwords.

MONDAY MORNING

Protect the accounts that can move money, reset access, reach many systems or authorize exceptions. Strong authentication is a workflow priority, not a company-wide checkbox.

CISA says businesses should aim for phishing-resistant MFA and start with administrative, remote-access and sensitive-data roles.

PATHWAY 02

The attack moved off email.

Email remains important, but voice, text, social media and collaboration tools create interactive pressure that traditional email training and filtering cannot see.

16%

Social Engineering breaches

Third-most common breach pattern.

[V26]

41%

Used a non-email vector

Within Social Engineering breaches.

[V26]

+40%

Higher median interaction

Phone-centric vs email simulations; n=35 vs 8,395.

[V26]

THE WIDENING CHANNEL MAP


EMAIL

links, files, callback


TEXT

smishing, QR, urgency


VOICE

help desk, executive


SOCIAL

recruiting, rapport


CHAT

Teams, support, access

DESIGN THE RESPONSE

- ✓ Teach verification for calls, texts and collaboration tools - not email cues alone.
- ✓ Give the help desk a high-risk reset script and a safe refusal path.
- ✓ Measure report rate and time-to-report across channels, normalized for scenario difficulty.

EVIDENCE LIMIT

The 40% comparison is simulation telemetry, not a real-world breach-success rate. Non-email data came from only 35 campaigns. Treat it as a directional warning, not a universal probability.

PATHWAY 03

The expensive attack is a believable request.

BEC can use spoofed mail, a compromised account, malware, voice or a deepfake. The target is not the inbox. It is the payment, payroll or vendor-change workflow.

\$3.05B

Reported BEC loss

24,768 IC3 complaints in 2025.

[F25]

\$1.164B

Attempted FFKC theft

Across 3,900 eligible incidents.

[F25]

\$679M

Funds frozen

58% of attempted dollars in the FFKC set.

[F25]

THE BEC KILL CHAIN

REQUEST

urgent or familiar

PAUSE

do not reply

VERIFY

known second channel

APPROVE

dual authority

RESPOND

bank + IC3 fast

SPEED IS A BUSINESS CONTROL

If a transfer is fraudulent, call the financial institution immediately, request recall/freezing action and file a detailed IC3 complaint. The 58% figure applies to eligible FFKC attempted dollars, not to every reported transfer.

MINIMUM GUARDRAIL

Any change to banking details, payroll, payment destination or gift-card request is verified using a known number or in person - never contact details supplied in the request.

PATHWAY 04

AI scales familiar attacks - and creates a new data path.

Current evidence shows AI accelerating impersonation, malware and reconnaissance while ordinary workplace use moves source code and sensitive documents into services the organization may not govern.

1 in 4

AI-enabled malicious breaches

IBM sample: 602 already-breached organizations.

[I26R]

\$893M

AI-described IC3 losses

22,364 complaints; descriptor may overlap categories.

[F25]

45%

Regular corporate-device AI users

Up from 15% in monitored telemetry.

[V26]

SHADOW AI SIGNAL

67%

of observed AI users used non-corporate accounts on corporate devices.

Source code was the most common submitted content. Research or technical documents appeared in 3.2% of GenAI-targeting DLP events. This is monitored-customer telemetry, not a workforce survey.

COST CONTEXT

\$6.0M

average AI-enabled malicious breach cost in IBM's studied sample, versus a \$4.99M global mean. This is not an SMB forecast and does not prove AI caused the difference.

CONTROL THE USE, NOT ONLY THE TOOL

APPROVE

named services and use cases

CLASSIFY

what data may be entered

ATTRIBUTE

human owner for every agent

LIMIT

least privilege and scoped tokens

LOG

prompts, outputs, actions, exceptions

PATHWAY 05

Your attack surface includes everyone connected to you.

Third-party access, unmanaged devices, exposed edge systems and credentials can turn an opportunistic intrusion into a ransomware or data-loss event.

48%

Third-party involvement

Across all confirmed breaches; up 60%.

[V26]

48%

Ransomware involvement

Across all confirmed breaches; up from 44%.

[V26]

96%

Known-size ransomware victims

SMBs in Verizon's size-known subset; sampling caveat.

[V26]

SMB ENTRY CONDITIONS IN VERIZON'S ANALYSIS

38%

compromised credentials

29%

unpatched edge devices

55%

third-party involvement

45%

human element present

These are different metrics with different denominators and may overlap. They are not parts of one pie.

SUPPLIER CONTROL PLANE

- ✓ Require MFA and named accounts for vendor access; remove standing privilege.
- ✓ Review remote tools, shared accounts and cloud permissions with the supplier.
- ✓ Contract for notification, evidence access and tested recovery responsibilities.

RESILIENCE TEST


Patch exposed edge systems, maintain isolated recoverable backups and rehearse identity containment. In Verizon's ransomware contributor subset, 69% of victims did not pay - a sign that resilience changes leverage.

WHAT CHANGES BEHAVIOR

Measure what people do - not what they watched.

Completion proves attendance. A human-risk program should also measure protective action, reporting speed, retained behavior and whether workflow controls stopped impact.

KNOWBE4 2026 PHISH-PRONE PERCENTAGE



42M simulations | 14.8M users | 64,000 organizations. Vendor customer telemetry; not randomized evidence of breach reduction.

USE THE TRAJECTORY CAREFULLY

The absolute percentages show strong association with an ongoing program. They do not establish how much training alone changed real-world compromise risk.

The publisher also labels the 33.2% to 4.2% change as both 79% and 87% in different 2026 materials. Raw values are shown here.

DATA INTEGRITY NOTE

OUTCOME LADDER



NIST recommends a life-cycle learning program with evaluation that progresses beyond reaction and knowledge toward behavior and organizational results.

HRS INTERPRETATION

Risk is not a person. It is a moment.

The same action can carry very different risk depending on the channel, role, privilege, workflow and recovery options. Measure the context, not only the click.


PRIORITIZATION, NOT PUNISHMENT

A useful executive index should show its drivers, direction, recency and data coverage. It should include protective behavior and recovery - not only failures. It should never be the sole basis for ranking people or making employment decisions.

CONCEPTUAL PRIORITY

Exposure x consequence x control gap x recency, reduced by guardrails and recovery capability. This is a decision aid, not a breach probability.

THE HRS OPERATING LOOP

One loop. Visible evidence. Continuous change.

A defensible human-risk program does not end with a course. It connects observed signals to a prioritized intervention and then checks whether the behavior, control or outcome changed.



CONNECT TO THE BUSINESS LIFECYCLE

GOVERN

IDENTIFY

PROTECT

DETECT

RESPOND

RECOVER

NIST CSF 2.0 keeps governance, protection, detection, response and recovery in the same conversation. Human risk should do the same.

Maturity Benchmark

Move from checkbox to adaptive.

Score each dimension separately. A polished training calendar should not hide weak payment controls, poor reporting culture or missing identity coverage.

1

CHECKBOX

Annual generic training. Completion is the only metric. No clear owner or report route.

2

BASELINE

Owner, policy, reporting path, MFA inventory and initial measurement exist.

3

CONTINUOUS

Role-based cadence, multi-channel exercises, verification and monthly review operate.

4

RISK-LED

Signals are contextualized. Interventions target repeat risk, impact and control gaps.

5

ADAPTIVE

Event-triggered coaching, control automation, privacy review and incident learning are embedded.

SCORE THESE DIMENSIONS INDEPENDENTLY

GOVERNANCE

WORKFLOWS

SIGNALS

GUARDRAILS

REPORTING


OUTCOMES

EXECUTIVE DASHBOARD

Show risk, control and recovery in one view.

The board needs direction, coverage, high-impact workflows and outcomes. Operators need denominators, freshness and the event trail underneath every number.

ILLUSTRATIVE STRUCTURE - NOT CLIENT DATA



MEASUREMENT PLAYBOOK

Every metric needs a decision attached.

Use a small, stable set that connects observed behavior to a control or outcome. Normalize by opportunity so a busy role is not mislabeled simply because it receives more attacks.

MEASURE	DEFINITION	DENOMINATOR	DECISION
Protective-action rate	Safe verification or report	Eligible opportunities	Reinforce
Time-to-report	Median and 90th percentile	Valid reports	Triage
PR-MFA coverage	FIDO2/passkey protected	High-impact accounts	Protect
Repeat-risk rate	Comparable repeat in 30/90d	Prior participants	Target
Workflow coverage	Second channel + dual approval	High-risk workflows	Control
Outcome impact	Verified loss, downtime, recovery	Material events	Govern

MEASUREMENT RULES

Difficulty-label simulations with the NIST Phish Scale. Segment by role and channel. Keep privacy, fairness and purpose limits explicit. Pair every exposure or failure metric with a protective-action or recovery metric.

A REPORT SHOULD CHANGE MONDAY MORNING

Build the control loop in 30 days.

The goal is not program perfection. It is visible ownership, protected high-impact identities, safer money workflows, one-click reporting and a baseline you can improve.

WEEK 1**Name the owner**

Assign one accountable executive and map the five workflows that can move money, expose data or grant access.

WEEK 1**Publish the report route**

One-click reporting, visible no-blame language and a bank/MSP/insurer/counsel/IC3 contact tree.

WEEK 2**Protect critical identities**

Phishing-resistant MFA for admin, email, banking, remote access and sensitive roles.

WEEK 2**Lock payment changes**

Known second-channel verification and dual approval for vendor, payroll and destination changes.

WEEK 3**Inventory exposure**

Supplier access, shared accounts, remote tools, BYOD, exposed credentials and workplace AI.

WEEK 3**Baseline behavior**

Difficulty-labeled email plus one voice, text or collaboration scenario; measure reporting too.

WEEK 4**Tabletop one handoff**

Run a supplier-payment impersonation or deepfake scenario from request through recovery.

WEEK 4**Review evidence**

Set the dashboard, owners, next interventions and a date for the 90-day review.

EVIDENCE OF COMPLETION

Named owners | MFA coverage export | payment-verification procedure | report-channel test | incident contact tree | baseline with cohort, difficulty and denominator | tabletop actions

DAYS 31 TO 365

Make the loop continuous.

After the first controls are in place, expand channel coverage, integrate signals, test recovery and prove whether repeat risk, reporting and material outcomes changed.

DAYS 31-90

Contextualize the baseline

- ✓ Segment learning by role, workflow and observed exposure.
- ✓ Run email, QR, text, voice, BEC and credential scenarios with difficulty labels.
- ✓ Give immediate supportive coaching and close the loop with reporters.
- ✓ Inventory AI use, non-human identities, supplier access and exceptions.
- ✓ Tabletop a supplier-payment deepfake or account-takeover event.
- ✓ Review operating metrics monthly and leadership outcomes quarterly.

MONTHS 4-12

Integrate and prove

- ✓ Expand passkeys/FIDO2 and remove standing privilege.
- ✓ Combine identity, exposure, phishing, triage, policy and incident signals.
- ✓ Automate event-triggered reinforcement with human approval.
- ✓ Calibrate scoring, time decay, privacy and fairness review.
- ✓ Extend measurement to vendors, contractors and AI agents.
- ✓ Reassess maturity quarterly and publish verified outcome changes.

PROOF STANDARD

Show the baseline, denominator, data coverage and intervention. Then show the comparable result, uncertainty and business outcome. If the evidence cannot support causation, say association.

METHODOLOGY

Make every headline auditable.

This report was rebuilt independently from current original sources. The attached report was used only to understand intended audience, scope and visual direction; its claims were rechecked rather than copied.

EVERY STATISTIC SHOULD SHOW

- ✓ Publisher, edition and publication date
- ✓ Observation period and geography
- ✓ Population, sample and denominator
- ✓ Exact definition and unit
- ✓ Known limitation or uncertainty
- ✓ Stable original-source link
- ✓ Retrieval date and version note

CLAIMS THIS REPORT AVOIDS

- ✗ 62% were caused by employees.
- ✗ AI caused \$893M in losses.
- ✗ Voice phishing is always 40% more successful.
- ✗ MFA stops 99% of attacks.
- ✗ A global mean predicts one SMB's loss.

DEFINITIONS
HUMAN ELEMENT

DBIR grouping that includes social engineering, error and misuse. Present does not mean caused.

SMB

Verizon defines small organizations as fewer than 1,000 employees; HRS's practical audience may be smaller.

BEC

Impersonation or account compromise used to redirect money or sensitive action; may involve malware.

PHISHING-RESISTANT MFA

FIDO2/WebAuthn security keys or passkeys; not every authenticator app or SMS code.

SHADOW AI

Use of AI services, accounts or extensions outside approved governance and monitoring.

PROTECTIVE ACTION

A safe verification, report, refusal, escalation or control use tied to a real opportunity.

General information only. This report is not legal, financial, insurance or incident-response advice. Regulatory scope and notification duties require fact-specific review.

SOURCE REGISTER

Direct links. Named limits. No hidden footnotes.

All sources were checked for this rebuild on 24 August 2026. Stable direct report links are used where available.

[V26] VERIZON 2026 DBIR

Published 19 May 2026. More than 31,000 incidents and more than 22,000 confirmed breaches in 145 countries; overall breach n=22,624; data 1 Nov 2024-31 Oct 2025.

Limit: Nonrandom contributor sample; categories overlap; year mix can change.

[OPEN PRIMARY SOURCE - verizon.com](#)

[F25] FBI IC3 2025

Released April 2026. 1,008,597 complaints and \$20.877B in reported loss; complaint, BEC, AI-descriptor and FFKC figures used here.

Limit: Underreporting, overlap and verification limits; reported loss is not total harm.

[OPEN PRIMARY SOURCE - ic3.gov](#)

[I26R] IBM COST OF A DATA BREACH 2026

Published 29 July 2026. Ponemon studied 602 breached organizations worldwide; breach period March 2025-February 2026.

Limit: Sponsored interview study of already-breached organizations; averages are not SMB forecasts.

[OPEN PRIMARY SOURCE - newsroom.ibm.com](#)

[K26] KNOWBE4 2026 BENCHMARK

Published 7 July 2026. 42M simulations across 14.8M users at 64,000 customer organizations; raw PPP values used.

Limit: Vendor customer data, not randomized. Publisher materials conflict on 79% vs 87% reduction.

[OPEN PRIMARY SOURCE - knowbe4.com](#)

[G26] GOOGLE M-TRENDS 2026

Published 23 March 2026. Mandiant analysis based on more than 500,000 incident-response hours during 2025.

Limit: Consulting cases are serious-event samples, not a random global prevalence survey.

[OPEN PRIMARY SOURCE - cloud.google.com](#)

[MS25] MICROSOFT DIGITAL DEFENSE 2025

Published 16 October 2025. Microsoft security telemetry for July 2024-June 2025; identity trends used as supporting signal.

Limit: Sign-in attacks are not confirmed breaches; product-population telemetry.

[OPEN PRIMARY SOURCE - microsoft.com](#)

[P25] PROOFPOINT SIMULATION ANALYSIS

Published 8 April 2025. More than 55,000 campaigns and 212M simulated messages sent during 2024.

Limit: Customer simulation data; template mix and program maturity affect results.

[OPEN PRIMARY SOURCE - proofpoint.com](#)

[N50] NIST SP 800-50 REV. 1

September 2024 guidance for a life-cycle cybersecurity and privacy learning program, behavior change, metrics and evaluation.

Limit: Normative guidance, not outcome prevalence data.

[OPEN PRIMARY SOURCE - doi.org](#)

[NCSF] NIST CSF 2.0 SMB GUIDE

February 2024 quick-start guide for SMB cyber-risk management across Govern, Identify, Protect, Detect, Respond and Recover.

Limit: Guidance must be tailored to business context and risk tolerance.

[OPEN PRIMARY SOURCE - doi.org](#)

[NPS] NIST PHISH SCALE USER GUIDE

NIST method for estimating human phishing detection difficulty so simulation results can be interpreted more fairly.

Limit: A difficulty model, not a predictor of real-world compromise.

[OPEN PRIMARY SOURCE - tsapps.nist.gov](#)

[CISA] CISA REQUIRE MFA

SMB guidance recommending phishing-resistant MFA and prioritization of admin, remote access and sensitive-data roles.

Limit: Implementation guidance; MFA types have materially different resistance.

[OPEN PRIMARY SOURCE - cisa.gov](#)

[FBEC] FBI BEC GUIDANCE

FBI guidance to verify account or payment-procedure changes using independently sourced contact details and report fraud quickly.

Limit: General protective guidance; recovery depends on timing and transaction facts.

[OPEN PRIMARY SOURCE - fbi.gov](#)

[NGAI] NIST AI 600-1

July 2024 Generative AI Profile with suggested actions to Govern, Map, Measure and Manage GenAI risk.

Limit: Voluntary cross-sector profile; select controls based on use case and risk.

[OPEN PRIMARY SOURCE - nvlpubs.nist.gov](#)



HumanRiskShield

HUMAN RISK MANAGEMENT

Measure the behavior. Fix the workflow. Prove the change.

HumanRisk Shield helps SMBs and the MSPs that protect them turn scattered human-risk signals into a continuous, evidence-led program.

01 IDENTIFY

priority roles, workflows and exposure

02 MEASURE

behavior, identity, reporting and controls

03 IMPROVE

targeted coaching plus practical guardrails

04 PROVE

direction, coverage, evidence and outcomes

START WITH THE HANDOFF THAT MATTERS MOST

HumanRiskShield.com

Research first. No fabricated numbers. Coaching without shame. Evidence you can take into an MSP review, leadership meeting or control audit.