THE COMMAND KIT · NO. 1 OF 2

The Top 10 Avoidable Incidents

A field guide to the person-patterns behind the breaches you can still prevent — and the early signals already sitting in your logs.

From the keynote **"Start With the Why"** by **Ashley Rose**, CEO & Co-Founder, Living Security
Pioneer of the Human Risk Management category · Forrester Wave Leader, 5.0/5.0 in Human Risk Quantification

livingsecurity

livingsecurity.com/command-kit

Every breach has two beginnings.

There's the date on the incident report. And there's the date — weeks earlier — when a password got reused, a reset got granted without verification, a consent got clicked, an export started climbing above baseline. Security has only ever told time from the first date. The second one was there all along, in data you already own.

Reading that second clock matters more now than it ever has, because the first clock is no longer winnable. Attackers now hand off initial access inside compromised environments in as little as **22 seconds** — down from more than eight hours in 2022.[1] No human-speed response wins that race. The race didn't get harder. It ended. The only clock left to win is the one that starts weeks before the incident report.

22s

Attacker hand-off time after initial access — down from 8+ hours in 2022.

Mandiant M-Trends 2026

10%

of the workforce drives roughly **73%** of risky behavior — validated across 100+ enterprises.

Living Security × Cyentia Institute, 2025

~1.5%

Median phishing-simulation click rate — a floor that years of blanket training haven't moved.

Verizon DBIR

That's why this guide names **people patterns, not malware signatures**. Every incident on this list forms the same way: a small, findable group of people (and increasingly, the AI agents acting on their credentials) behaving their way toward a specific breach — visibly, over weeks, in signals your identity, email, endpoint, SaaS, and HR systems already capture. The problem you thought was everyone is actually someone. And someone is findable.

> **Your workforce isn't the weakest link.**
> **It's the largest sensor network you never turned on.**

How to use this guide. Don't adopt all ten. Pick the one that maps to the breach your organization refuses to accept — the one your SOC is already tired of responding to. Each entry gives you the pattern, the second-clock signals that show up before it lands, the evidence it's real, and the two ways to shrink it: **engage the workforce** (targeted training, testing, nudges, scorecards) or **protect the environment** (access restriction and orchestrated controls). Then run the companion piece in this kit — *The Named-Breach Working Session* — against the one you picked.

Nothing in this guide requires a purchase to go look for. You already own every signal on these pages.

1 "The Backup Keyholders"

BECOMES: RANSOMWARE-DRIVEN INTRUSION

The employees and admins whose credentials and behavior stand between an intruder and your core systems — domain admins, storage and backup operators, and the everyday users whose reused or already-stolen passwords open the initial path. Ransomware doesn't start with encryption; it starts with one of these people, weeks earlier.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Passwords reused across systems by privileged users
- Employee credentials surfacing in infostealer logs and breach dumps
- MFA-fatigue approvals on privileged accounts
- Backup and storage consoles reachable with standing credentials

THE PLAY

- **Engage:** targeted credential-hygiene training and nudges for the keyholder cohort — at the reuse moment, not annually
- **Protect:** rotate exposed credentials, phishing-resistant MFA on privileged access, remove standing paths to backup infrastructure

Evidence: Ransomware appeared in **48% of confirmed breaches** — the highest share in the report's history, up from 44% the year before (Verizon 2026 DBIR).[2]

2 "The Gatekeepers"

BECOMES: HELP-DESK SOCIAL ENGINEERING & ACCOUNT TAKEOVER

The help-desk and support staff with the authority to reset passwords and MFA. The Scattered Spider playbook — the one behind the 2025 UK retail attacks and the wave that hit airlines — is built entirely on talking a Gatekeeper into a reset. AI voice cloning has made the impersonation dramatically more convincing, and the caller answers every verification question correctly.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Resets granted without full verification
- Reset-velocity outliers in the ticket queue
- Lookalike "sso" / "helpdesk" domains registered against your brand
- Verification steps skipped when the caller creates urgency

THE PLAY

- **Engage:** voice-phishing simulation against the help desk itself; train the "no-exception" callback protocol under pressure
- **Protect:** step-up verification for MFA resets, restrict who can reset privileged accounts, monitor lookalike-domain registrations

Evidence: In Verizon's 2026 DBIR phishing-simulation data, voice- and text-based lures drew median click rates roughly **40% higher than email**. Scattered Spider's 2025 help-desk campaigns disrupted Marks & Spencer, Co-op, and Harrods, then WestJet, Hawaiian Airlines, and Qantas.[3]

3 "The Patch Deferrers"

BECOMES: EXPLOITATION OF A KNOWN, UNPATCHED VULNERABILITY

The system owners and admins who keep clicking "remind me later." Each deferral is rational in the moment — a release window, a fragile dependency, a busy week. Strung together, deferrals become the open door attackers now favor most. This is a behavior pattern wearing a vulnerability-management costume.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Patch windows deferred again — and again — by the same owners
- Maintenance tickets aging past SLA on internet-facing systems
- Exception requests renewed without re-review
- The same assets missing every scan cycle

THE PLAY

- **Engage:** surface deferral streaks to the owners and their leaders via scorecards — risk in front of the people who own it
- **Protect:** auto-isolate or compensating controls for systems past deferral thresholds; make "defer" cost a decision, not a click

Evidence: In the Verizon 2026 DBIR, **vulnerability exploitation overtook credential theft** as the top initial breach vector.[4]

4 "The Shadow Adopters"

BECOMES: SENSITIVE DATA LEAKING THROUGH UNMANAGED AI & APPS

The employees quietly adopting AI tools and apps your security team has never reviewed — usually to be more productive, almost never to do harm. They paste the strategy doc, the customer list, the source code. No malware. No intrusion. Nothing for detect-and-respond to detect — just data leaving through a personal login.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Personal-account logins to AI platforms from corporate devices
- Uploads and pastes into unmanaged tools trending up
- Consents granted to unvetted browser extensions and apps
- Sanctioned-tool adoption lagging in specific teams

THE PLAY

- **Engage:** nudge at the paste-and-upload moment; route the cohort to the sanctioned alternative instead of punishing the productivity
- **Protect:** block personal-account access to high-risk tools; give the workforce an approved AI path so the shadow one isn't worth it

Evidence: 67% of enterprise AI activity runs through personal accounts — often on the very platforms the company licenses (Teramind, 2026). Source code consistently ranks among the top categories of confidential data flowing into external AI tools (Cyberhaven).[5]

5 "The Authorized Approvers"

BECOMES: BUSINESS EMAIL COMPROMISE & PAYMENT FRAUD

The people with the authority to move money — AP teams, finance leads, executive assistants, deal owners. BEC doesn't break anything technical: the money moves because a real human approves it inside a real workflow, convinced by a message that looked exactly right. AI-written lures have removed the tells everyone was trained to spot.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Banking-detail changes accepted without out-of-band callback
- First-time payee approvals climbing
- Payment approvals processed under time pressure, outside hours
- Lookalike supplier and executive domains registered

THE PLAY

- **Engage:** simulate executive- and vendor-impersonation against the approver cohort specifically; rehearse the verification habit
- **Protect:** enforced dual approval and out-of-band verification for payee changes above threshold — no exceptions for urgency

Evidence: Reported cybercrime losses reached **$20.9 billion in 2025**, up 26% year over year, with business email compromise among the costliest categories (FBI IC3 2025 Internet Crime Report).[6]

6 "The Recyclers"

BECOMES: CREDENTIAL-STUFFING & ACCOUNT TAKEOVER THAT FEEDS RANSOMWARE

The employees who reuse one password across work and personal life. When any of those services is breached — or an infostealer sweeps a personal laptop — the corporate credential enters the criminal supply chain, where it's validated, sold, and used. This is the second clock at its most literal: the attack is purchased months before it happens.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Passwords reused across systems (visible at reset and SSO)
- Corporate emails appearing in fresh breach corpora and stealer logs
- Logins from devices with infostealer indicators
- Impossible-travel and off-pattern authentications that "succeed"

THE PLAY

- **Engage:** nudge at the reuse moment; password-manager adoption targeted at the recycler cohort, not the whole company
- **Protect:** auto-rotate exposed credentials, monitor stealer-log exposure, require step-up auth for flagged accounts

Evidence: A majority of ransomware victims — **54% and rising** — had employee credentials circulating in infostealer logs before the attack, and half of prior infections occurred within roughly **95 days** of the ransomware event: a three-month head start, on sale (SpyCloud).[7]

7 "The Power Connectors"

BECOMES: OVER-SCOPED THIRD-PARTY & OAUTH COMPROMISE

The integrators and early adopters who wire apps together — granting OAuth scopes, installing connectors, authorizing tokens that then act with machine speed and standing access. Every consent is a new non-human identity in your environment. When the vendor behind one of those tokens is compromised, the attacker inherits everything your people approved — and MFA never gets asked.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Consents to over-scoped apps (read-everything, act-as-user)
- Dormant integrations holding live, long-lived tokens
- Non-human identities multiplying with no named owner
- Token activity outside the app's normal purpose or hours

THE PLAY

- **Engage:** make the connector cohort review their own grants — "here's what this app can still do; do you still need it?"
- **Protect:** least-scope consent policies, expiry on tokens, revoke dormant grants, alert on scope escalation

Evidence: In August 2025, attackers who compromised the Salesloft Drift integration used its stolen OAuth tokens to export data from the Salesforce environments of **700+ organizations** — bypassing MFA entirely, through access humans had legitimately authorized.[8]

8 "The Expedient Fixers"

BECOMES: EXPOSURE THROUGH WORKAROUNDS & RISKY SHORTCUTS

The high performers who route around friction to get work done — sharing credentials in a ticket, disabling the agent that slows the build, moving files through a personal drive, leaving the firewall exception open "for now." They aren't malicious; they're productive. But the workaround that works today is the finding in tomorrow's post-mortem, and no alert fires when it happens.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Security tools disabled or paused on endpoints
- Credentials shared in tickets, chat, and code
- Personal cloud and file-sharing use climbing in one team
- The same control exception requested repeatedly

THE PLAY

- **Engage:** treat repeat workarounds as a signal of broken process — fix the friction, then nudge the cohort to the paved path
- **Protect:** tamper protection on controls, secrets scanning where credentials get shared, sanctioned fast lanes for legitimate urgency

Evidence: This is the pattern blanket training can't reach: median phishing-simulation click rates have sat at a **~1.5% floor** for years of industry-wide training (Verizon DBIR) — because awareness isn't the problem. Friction is. The fix is behavioral and environmental, not another course.[9]

9 "The Quiet Exit"

BECOMES: INSIDER DATA MISHANDLING BEFORE DEPARTURE

The employee who has already decided to leave — sometimes before anyone else knows. The customer list, the pipeline export, the design files: they walk out through downloads and forwards in the weeks between the decision and the goodbye. No malware. No intrusion. Nothing for detect-and-respond to detect. The only fingerprint is behavior against baseline.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Exports climbing above the person's own baseline
- Bulk downloads and repository access outside their role
- Forwarding to personal email; syncs to personal devices
- HR signals — resignation filed, performance cycle, role change

THE PLAY

- **Engage:** a manager conversation at the first off-baseline signal — most walk-outs stop when someone simply asks
- **Protect:** join HR and security signals so departure windows get scrutiny by default; step-down access at notice, not at exit

Evidence: This is the canonical second-clock incident: every signal above predates the loss, and every one lives in systems you already run — SaaS audit logs, DLP, email, HRIS. The gap is that nobody joins them to a person in time.

10 "The Standing Privileged"

BECOMES: PRIVILEGE ABUSE & BLAST-RADIUS AMPLIFICATION

The people — and service accounts, and now AI agents — holding admin rights they no longer use or never needed. Standing privilege is quiet risk: it does nothing wrong until any other incident on this list touches one of these accounts, and then it turns a contained event into an enterprise one. Identity governance sees who has access. It doesn't see that the access hasn't been used in 14 months.

THE SECOND CLOCK — ALREADY IN YOUR LOGS

- Admin entitlements unused for months but never revoked
- Privileges surviving role changes and reorgs
- Service accounts and agents with no named owner
- MFA-fatigue approvals on high-privilege identities

THE PLAY

- **Engage:** put unused-privilege scorecards in front of the owners; make keeping an entitlement an active choice
- **Protect:** just-in-time elevation instead of standing rights; auto-expire on role change; owner-of-record required for every non-human identity

Evidence: The workforce now includes agents acting on enterprise credentials at machine speed. The 2025 Drift incident showed what standing, over-scoped access does at scale: tokens with broad rights, no expiry, and no behavioral oversight became a 700-organization breach.[8]

You don't need ten enemies. You need one.

Every pattern in this guide forms weeks before it lands, in signals you already own. Naming yours isn't fear — it's focus. So don't circulate this list and move on. Pick the one incident your organization refuses to accept, and aim everything at it.

THE OPERATING LOOP

Name it — the breach you refuse to accept → **Find the few** — the cohort driving it, in your data → **Act** — engage people or adapt the environment → **Prove it** — the number set up front, reported ↻

Worked example — "The Recyclers": stolen credentials named as the breach → the reuse cohort identified → rotate + nudge at the reuse moment → exposed-credential count, down and reported to the board. The number moved or it didn't. That's the whole ROI conversation.

Three questions to take into your next leadership meeting:

THEY ASK TODAY	ASK INSTEAD
"What's our risk posture?"	**"What breach do we refuse to accept?"**
"Who completed the training?"	**"Who's trending toward that breach?"**
"What does the tool cost?"	**"What number proves it isn't forming?"**

Then run the companion piece in this kit — **The Named-Breach Working Session** — a 60-minute agenda that takes your team from "interest in human risk" to a named breach, an identified cohort, and a metric your CFO will fund. None of it requires buying anything.

SOURCES

1. Mandiant, *M-Trends 2026* — median time from initial access to attacker hand-off: 22 seconds, down from 8+ hours in 2022.
2. Verizon, *2026 Data Breach Investigations Report* — ransomware present in 48% of confirmed breaches, the highest share in report history.
3. Verizon, *2026 DBIR* — phishing-simulation median click rates for voice/text lures ~40% higher than email. Scattered Spider 2025 campaigns: public reporting on Marks & Spencer, Co-op, Harrods, WestJet, Hawaiian Airlines, Qantas.
4. Verizon, *2026 DBIR* — vulnerability exploitation surpassed credential abuse as the leading initial breach vector.
5. Teramind, *Shadow AI Behavior Report* (2026); Cyberhaven AI usage research on confidential data categories flowing to AI tools.
6. FBI Internet Crime Complaint Center, *2025 Internet Crime Report* — $20.9B in reported losses, +26% year over year.
7. SpyCloud ransomware and identity-exposure research — 54%+ of ransomware victims with prior credential exposure in infostealer logs; 50% of prior infections within ~95 days of the attack.
8. Public incident reporting on the August 2025 Salesloft Drift OAuth compromise affecting 700+ organizations' Salesforce environments.
9. Verizon DBIR — median phishing-simulation click rate plateau (~1.5%); recently trained employees report at ~4× the base rate.
10. Living Security × Cyentia Institute, *2025 State of Human Cyber Risk Report* — ~10% of users drive ~73% of risky behavior; behavioral data from 100+ enterprises.

Don't start with the tool. Start with the breach you refuse to accept.

Living Security pioneered Human Risk Management: unifying the behavior, access, and threat signals already in your stack to find the ~10% driving ~73% of workforce risk — human and AI agent alike — change the behavior, and prove the reduction. Forrester Wave Leader. 5.0/5.0 in Human Risk Quantification.

Get the working session, run it with your SOC, and if you can't find your 10% in your own data — you just found the gap. We should talk.