



THREAT INTELLIGENCE BRIEF

2026 Mid-Market Cyber Threat Trends Report

Three shifts in this year's breach data, and what they mean for lean IT and security teams.

Why Size Is Not a Defense

You’ve probably had this thought at some point: we’re not a bank, we’re not a hospital system, nobody’s targeting us specifically. It’s a reasonable thing to think. It’s also not how the last year of breach data reads.

Most attacks on small and mid-sized companies aren’t targeted at all. They’re opportunistic. Automated scanning finds an exposed system, a criminal group runs the same playbook it ran on the last thousand victims, and the fact that your company wasn’t specifically chosen doesn’t change what happens next. Of the ransomware victims in this year’s breach data where organization size was known, 96 percent were small or mid-sized businesses.


96%

Of 2026’s ransomware victims, where organization size was known, were small or mid-sized businesses.

VERIZON 2026 DATA BREACH INVESTIGATIONS REPORT

That’s not because small and mid-sized companies have more worth stealing. It’s because they’re easier to get into. The question was never whether your company is a big enough prize. It’s whether your company is an easy enough door.

The same data, broken down further, makes the same point with more specificity. Looking only at the actions involved in confirmed SMB breaches, ransomware appears in 83 percent of them, use of stolen credentials in 39 percent, and exploited vulnerabilities in 30 percent. The pattern points less to industry or revenue than to two ordinary conditions: compromised credentials, present in 38 percent of these cases, or an unpatched vulnerability in an edge device, present in 29 percent.



This report looks at three specific findings from this year’s major industry breach and claims data and what they mean if you’re running IT or security at a company with a few hundred to a couple thousand employees, without a dedicated SOC and without unlimited budget.

None of these three findings are exotic attack techniques. They’re all quiet failures of things most companies already believe they’ve handled: backups, patching, and funds-transfer controls. Each one is worth checking against your own environment before you find out the hard way whether the assumption held.

WHAT'S INSIDE

01

Backup Infrastructure Is a Primary Target, Not a Fallback

Recovery denial: attackers disable backups, identity, and virtualization before any encryption starts.

02

Patch Cycles Are Losing Ground to Faster Exploitation

Exploited vulnerabilities passed stolen credentials as the top way in, while remediation rates fell.

03

Funds Transfer Fraud No Longer Needs a Human Mistake

A growing share of losses runs straight to the bank, with no phishing email and no employee in the path.

Backup Infrastructure Is a Primary Target, Not a Fallback

“We have backups” used to be the end of the conversation about ransomware. It isn’t anymore.

This year’s data documents a shift the industry is now calling recovery denial. Ransomware operators are no longer treating encryption as the main event. Before they encrypt anything, they go after the systems an organization would use to recover: backup infrastructure, identity services like Active Directory, and the virtualization layer that runs most companies’ servers. The goal isn’t just to lock your data. It’s to make sure that paying the ransom is the only way to get it back.

14 days

Median attacker dwell time in 2025, up from 11 days the year before. That is more time to locate and disable recovery systems before acting.

MANDIANT M-TRENDS 2026

In practice, this looks like: attackers exploiting misconfigured Active Directory settings to create administrator accounts that skip multi-factor authentication entirely, deleting backup files directly from cloud storage, and copying virtual machines, including the very identity systems that would normally catch this activity, so the copies can be picked apart without tripping any alarms. Named ransomware operations, including groups known as Akira and Qilin, have been observed running this playbook specifically.

THE RECOVERY-DENIAL SEQUENCE

01

Initial access via exploited vulnerability



02

Pivot to identity systems (Active Directory, credential vaults)



03

Backup and virtualization infrastructure disabled or copied



04

Production systems encrypted, recovery path already gone

A simplified composite of the pattern described in M-Trends 2026, not a literal reproduction of any single incident.

What this means for your environment: if your IR plan or DR runbook was last tested against “the building burns down,” not “the thing you’d recover with is gone before you know you’ve been hit,” it’s testing the wrong scenario. The relevant question isn’t whether backups exist. It’s whether they’re reachable, current, and isolated enough that an attacker sitting inside your network for two weeks couldn’t find and disable them before you ever see a ransom note. For a mid-market company, backup posture has moved from an IT hygiene line item to a primary attack surface, whether or not budgeting, testing, or prioritization has caught up with that shift.

Patch Cycles Are Losing Ground to Faster Exploitation

For most of the last two decades, stolen credentials, not unpatched software, were how attackers got in the door. That changed this year.

31%

Of initial access in 2026's confirmed breaches came through exploited vulnerabilities, up from 20% the year before. That is the first time exploitation has outpaced stolen credentials in the DBIR's 19-year history.

VERIZON 2026 DATA BREACH INVESTIGATIONS REPORT


Vulnerability exploitation rose from 20 percent of initial access vectors to 31 percent in a single year, a 55 percent increase. Credential abuse, the long-standing top vector, dropped to 13 percent, though part of that drop is a methodology change rather than a pure decline: this year's report split out pretexting, a more involved, real-time form of social engineering, as its own tracked vector, where it had previously overlapped with credential abuse counts. On a basis comparable to last year, credential abuse would sit closer to 16 percent, a real decline, just a smaller one than the headline number suggests.

Remediation moved in the opposite direction. Of the vulnerabilities on CISA's Known Exploited Vulnerabilities list, organizations fully remediated only 26 percent in 2025, down from 38 percent the year before. Median time to remediate rose from 32 days to 43. At the same time, the median number of KEV vulnerabilities an organization had to patch rose to 16 in 2025, up from 11 the year before, nearly 50 percent more in a single year.

INITIAL ACCESS VECTOR



KEV REMEDIATION RATE



What this means for your environment: a patch cadence built around monthly or quarterly cycles is now competing against exploitation windows measured in days, a pace that cadence was never built to match. This isn't a call to patch everything immediately. It's a call to know which handful of exposed, internet-facing systems would actually get you breached if left unpatched for six weeks, and to treat those differently than the rest of the queue. Most lean IT teams don't have that list today. That's a gap worth closing before it closes itself for you.

Funds Transfer Fraud No Longer Needs a Human Mistake

Most defenses against funds-transfer fraud assume a way in: a convincing email, a spoofed invoice, a moment where someone gets fooled. In 2025, a meaningful share of it stopped needing that moment at all.

39%

Of 2025 funds-transfer-fraud events involved no confirmed email compromise at all.

COALITION 2026 CYBER CLAIMS REPORT

A fifth of all funds-transfer-fraud events involved fraudulent instructions sent directly to a bank, bypassing any employee interaction entirely. Rather than spending time manipulating someone into acting, attackers are increasingly using stolen credentials to reach financial systems directly, a faster and quieter route to the same payout.

52%

of funds-transfer fraud still originates as business email compromise

39%

involved no confirmed email compromise at all

20%

went directly to fraudulent bank instructions, bypassing employees entirely

These are different cuts of the same claims data, not mutually exclusive shares of a whole.

Email-based deception hasn't gone away. It's still the largest single starting point. But a growing share of losses is happening on a route that doesn't leave "an employee who got fooled" as the point of failure, and doesn't leave a phishing email as the piece of evidence to point to afterward.

What this means for your environment: funds-transfer-fraud controls built entirely around training people to spot a suspicious email are addressing a shrinking share of how the fraud actually happens. If out-of-band verification for wire changes still depends on someone noticing a bad email first, it's not covering the 39 percent that never produces one. That control needs to trigger on the transfer itself, not on a suspicious message.

What This Means for Your Environment

None of the three trends in this report describe a novel attack technique or an exotic new threat. They describe the same set of systems most mid-market organizations already assume are handled: backups, patch management, and funds-transfer controls. That overlap with ordinary operations, not some exotic technique, is what makes them worth paying attention to.

Go back to where this report started. Being a mid-market company was never a reason attackers would skip you. It was a reason they'd find you first. The three shifts above are what finding you first looks like in practice this year: recovery infrastructure targeted before encryption even starts, exploitation outpacing patch cycles built for a slower era, and funds-transfer fraud increasingly bypassing the human interaction most controls are built to catch.

If any of these three raises a question about your own environment, that's worth a conversation, not a guess.

If you want to talk through what any of these three trends would mean for your own environment, we're glad to have that conversation.

CONTACT

Satine Technologies · Atlanta, Georgia
info@satinetech.com · satinetech.com



Sources: Verizon 2026 Data Breach Investigations Report; Mandiant M-Trends 2026; Coalition 2026 Cyber Claims Report. Figures are reproduced from those published reports; Satine Technologies is not affiliated with, and does not speak for, their publishers.

Glossary

Plain definitions for the terms used in this report, written for readers who own the risk without owning the tooling.

Active Directory	Microsoft's directory and identity service. It controls who can log in to what across a Windows network, which makes it a high-value target: control of it usually means control of everything else.
Business email compromise (BEC)	Fraud that begins with an attacker gaining access to, or convincingly imitating, a legitimate email account in order to redirect a payment or extract information.
CISA	The US Cybersecurity and Infrastructure Security Agency. Among other duties, it publishes the KEV catalog of vulnerabilities confirmed to be under active attack.
DBIR	The Verizon Data Breach Investigations Report, an annual study of confirmed breaches that many teams treat as the industry baseline for how attacks actually begin.
DR runbook	Disaster recovery runbook. The step-by-step procedure for restoring systems and data after an outage, whether caused by hardware failure, disaster, or attack.
Dwell time	How long an attacker is present inside an environment before being detected. Longer dwell time means more opportunity to locate backups, identity systems, and financial workflows.
Funds transfer fraud (FTF)	Any scheme that causes an organization to send money to an account the attacker controls, whether through a forged invoice, an altered wire instruction, or direct instructions to a bank.
Initial access vector	The specific method an attacker used to get inside in the first place, such as an exploited vulnerability, stolen credentials, or a phishing message.
IR plan	Incident response plan. The documented process for who does what once a security incident is detected, from the first alert through containment, eradication, and recovery.
KEV (Known Exploited Vulnerabilities)	A catalog maintained by CISA listing vulnerabilities confirmed to be under active exploitation. It is the practical short list for patch prioritization.
Out-of-band verification	Confirming a request through a separate channel from the one that carried it, such as calling a known phone number to verify a wire change that arrived by email.
Pretexting	Social engineering in which the attacker builds a believable scenario and interacts with a target in real time, rather than relying on a single deceptive message.
Ransomware	Malware that encrypts or removes access to an organization's data, after which the operator demands payment. Most operations now also steal data to create a second point of leverage.
Recovery denial	An approach in which attackers disable or destroy the means of recovery, such as backups, identity services, and virtualization infrastructure, before encrypting anything, so that paying is the only fast option.
Virtualization layer	The software that runs multiple virtual servers on shared hardware. Compromising it can expose or disable every system running on top of it at once.