SpyCloud

2026 Forecast Report

A year in review and look ahead to what we predict for 2027

Security and Identity teams are constantly tasked with shifting defense priorities.

Most security teams spend the first half of the year weighing where to look first. At the end of the year, we look back at what happened and start sizing up what to prioritize next.

It can feel like a constant game of catch-up. Plans for the new year are getting set right now, so we're not waiting for January to weigh in.

This year's report examines how 2026 played out, and grades how **last year's predictions ↗** fared. Then we'll put our experts' fresh takes for 2027 on the table.

The outlook draws on what we see in the criminal underground, our SpyCloud Labs research and monthly cybercrime updates, input from our team of experts, and over a decade of disrupting cybercrime as our mission. We also surveyed 50 security professionals at large enterprises about what they expect in 2027.

Table of Contents

1.	**Key takeaways**	**4**
2.	**Meet the experts**	**5**
3.	**Lookback: How did our previous predictions hold up?**	**6**
4.	**2026 in cybercrime updates**	**10**
	Pulse check: What teams dealt with in 2026	12
5.	**Outlook: What we expect to see next**	**13**
	Pulse check: What's on teams' minds for 2027	21
6.	**Putting it together**	**22**

Key takeaways

MFA is routinely bypassed:
48% of surveyed organizations experienced account takeovers despite having MFA enabled, primarily through session and device hijacking.

AI tokens are among the new targets:
Infostealers like WardenStealer are pivoting to collect API keys and tokens used by AI agents and coding tools.

Trusted access risks persist:
68% of enterprises were affected by vendor or SaaS breaches, highlighting that "trusted" integrations are now a primary attack vector.

Synthetic identities evolve:
Fraudsters are combining real stolen data with AI-generated media to create highly convincing synthetic identities.

Phishing targets non-human identities:
As organizations adopt AI agents with internal system access, attackers are beginning to phish these non-human identities (NHIs).

Meet the experts

We asked eight SpyCloud experts, from threat research to product to the people who work with defenders every day, what they're seeing and what security teams should have on their radar in 2027.



Damon Fleury

Chief Product Officer



Trevor Hilligoss

Chief Intelligence Officer & Head of SpyCloud Labs



Jason Lancaster

Chief Investigations Officer



Joe Roosen

Sr Director, Security Research



Andy Culler

Director of Security Research Engineering



Aurora Johnson

Director, Threat Research & Intelligence



Wallis Romzek

Director, Data Science



Donovan Stevens

Principal Sales Engineer

LOOKBACK

How did our predictions for 2026 hold up?

Rather than rehash the last 12 months, we're hitting the highlights. Spoiler: Six of last year's ten held up, and the other four are still developing.

1 ✓ CONFIRMED

The industrialization of cybercrime will further commoditize plug-and-play identity attacks.

Last year's hunch → Competition was making phishing kits cheaper and easier to use, which lowered the barrier to entry.

How it played out → Phishing kits now run like subscription products. **BlueKit bundles ↗** 80+ templates, text-message delivery and an AI assistant in one panel, and **Kali365 sold ↗** access to about 500 affiliates at $250 a month.

2 DEVELOPING

The threat actor community will keep splitting and shifting due to youth involvement, platform migrations, and global threat expansion.

Last year's hunch → **Telegram's crackdown ↗** was pushing criminals to platforms with looser moderation, and easy money kept drawing in teenagers.

How it played out → ShinyHunters ran data theft extortion, taking data from cloud accounts at **Udemy and Vimeo ↗** and breaching Canvas's parent company, **Instructure ↗**. SpyCloud parsed 6.38 billion unique **Chinese PII records ↗**, and it ingested 2.9x its usual yearly count of **Latin American breaches ↗** by July. **Handala leaked ↗** historical emails from the FBI director's personal account. Card traders landed on **Threads ↗**, but the updates don't show a move to X or WhatsApp, and forums kept coming back **under new names ↗**.

How did our predictions for 2026 hold up?

3 ✓ CONFIRMED

Identity (and non-human identity) sprawl will be tough to wrangle.

Last year's hunch → People had more exposed data than anyone had visibility into, and machine logins skipped the protections people get.

How it played out → Stolen machine logins drove the year's biggest incidents. The **Klue breach ↗** targeted Salesforce tokens, **TeamPCP used ↗** compromised developer credentials to slip malware into open-source packages, and **Remus ↗**, a new infostealer, copies the credential folders AI assistants use. In our June 2026 survey, **31% of affected organizations ↗** said exposed or overprivileged NHIs were their most common way in, nearly double phishing.

4 DEVELOPING

Insider threats will become mainstream.

Last year's hunch → Compromised accounts, employee cyber negligence and fraudulent new hires were blurring the lines between insider and interloper.

How it played out → Democratic People's Republic of Korea (DPRK) **fake-worker schemes ↗** are still prolific. The **FBI has busted laptop farms ↗** that fake workers used to proxy in, and SpyCloud's investigators have supported DPRK worker cases for multiple customers. In the Justice Department's case, **victim employers ↗** spanned health care, communications, cyber support, retail and manufacturing.

5 ✓ CONFIRMED

AI-enabled cybercrime & the AI arms race will play out – for bad or good.

Last year's hunch → Early AI use was mostly speeding up familiar attacks, and analysts saw room for more.

How it played out → AI phishing and AI coding both showed up. BlueKit has an AI assistant, **EvilTokens' ↗** chatbot could analyze up to 5,000 emails at once, and the **FortiBleed operators built ↗** tools with an AI code editor. Self-rewriting AI malware hasn't appeared. The twist is that AI errors cost criminals credibility, **including a retracted claim ↗** about 3.1TB of stolen data that an AI summary misread, and **a buyer's complaint ↗** about datasets of people who don't appear to exist.

How did our predictions for 2026 hold up?

6 ✓ CONFIRMED

Attackers will continue to innovate to bypass defenses.

Last year's hunch → Stolen sessions made MFA beatable, and malware was already slipping past endpoint tools.

How it played out → Each defense drew a new workaround. EvilTokens, live since February '18, made device code phishing commonplace, reaching at least 8,708 accounts across 6,585 corporate domains in 79 countries, according to SpyCloud's data. BlueKit enrolled its own authenticator in 38% of its compromises, keeping access alive after a password reset, and Remus blinds endpoint security tools before it runs.

7 ✓ CONFIRMED

Pathways to supply chain compromise will multiply.

Last year's hunch → Vendors and contractors held privileged access that few organizations audited.

How it played out → Victims were hit through someone else's systems again and again. Klue, which targeted Salesforce tokens, was the third such incident after Salesloft Drift and Gainsight. **TeamPCP ↗** compromised open-source packages including PyTorch Lightning, and five forum accounts posted 15 leaks from customers of one French SaaS platform, **BlgCloud ↗**. Taken together, the year's incidents show how many of the biggest compromises stemmed from a vendor or tool the victim already trusted.

8 DEVELOPING

Synthetic identity fraud will gain more momentum using stolen data and AI.

Last year's hunch → Leaked personal data was easy to piece into fake people, and AI was making them more convincing.

How it played out → A leaked dataset attributed to **IDMerit ↗** held about 122 million national IDs, and a **Chinese dataset ↗** held about 2.5 billion, which is the raw material for building fake people. SpyCloud Labs is also seeing criminals use AI to **generate fake datasets ↗** of people who don't exist and pass them off as real breaches. That isn't fraud in itself, but it shows how little effort it takes to produce believable identities. Many get called out as AI slop, but the low effort means they keep trying, and Labs expects them to continue growing and improving with each new iteration.

How did our predictions for 2026 hold up?

9 ✓ CONFIRMED

Distractions like megabreaches and combolists will keep popping up.

Last year's hunch → Criminals were repackaging old breach and infostealer data to make noise, and defenders kept chasing it.

How it played out → Big claims kept outrunning the facts. A Chinese dataset of 8.7 billion records, pieced together from earlier breaches, held 6.38 billion unique ones. A **fake ransomware group ↗** fooled security vendors before admitting a hoax, and unrelated actors reposted **ShinyHunters' data ↗** across forums for clout. Part of what sustains this trend is that some countries require organizations to report a suspected breach within days, before they've had time to confirm it. In **Japan ↗**, for example, a preliminary report is due within three to five days, so even a convincing fake can put a company under pressure.

10 DEVELOPING

Defenders, teams, and workflows will morph – and unite.

Last year's hunch → Security talent was scarce, so teams were expected to share work and lean on automation.

How it played out → Inside enterprises, the shift is still underway. Labs expects silos between SOC, threat analysis and fraud teams to give way as AI helps attackers move faster, and **35% of organizations ↗** plan to automate identity-related incident response. The clearest unity in 2026 happened between organizations. **The takedown of EvilTokens ↗** is a great example, with Microsoft leading the effort alongside law enforcement and private industry, and SpyCloud's data helping identify compromised accounts. SpyCloud Labs sees private companies increasingly bridging gaps between law enforcement and industry stakeholders by sharing data.

2026: The year in cybercrime updates

SpyCloud Labs ↗ – SpyCloud's in-house security research team – continuously monitors and analyzes the activities, and sometimes scuttlebutt, of the criminal underground to bring you a monthly cybercrime update covering emerging threats and trends. One standout is the drive from credentials to stolen sessions as a primary access vector. Plus the ongoing coordination of cybercrime disruption activities increases the cost to criminals.

2026 Cybercrime at-a-glance

A timeline of activity, research, and takedowns from the criminal underground this year

Takedowns and arrests

Breaches and extortion

New attacker tools

JAN

Malicious AI agent skills delivered Atomic Stealer, and stolen card data surfaced on Threads

FEB

Operation Red Card 2.0 made 651 arrests, and an exposed IDMerit database held about 875 million names

MAR

Tycoon 2FA and LeakBase fell, LeakBase's owner was arrested in Russia, and TeamPCP hit Cisco and the European Commission

APR

ShinyHunters took data from cloud accounts at Udemy and Vimeo, and SpyCloud parsed 6.38 billion unique Chinese PII records

MAY

ShinyHunters pushed ransom pop-ups to Canvas users during finals, and Kali365 faked its own shutdown

JUN

The Klue breach exposed Salesforce tokens, Operation Endgame hit StealC, and FortiBleed leaked 73,932 firewall logins

JUL

A fake ransomware group fooled security vendors, LeakBase returned, the Jalisco kit beat device code expiry, and Latin American breaches ran 2.9x the usual rate

AUG

Five forum accounts coordinated 15 leaks from one French SaaS platform

SEP

BlueKit introduced browser-in-the-middle phishing, EvilTokens fell, and ShinyHunters hijacked Cl0p's site as two of its members were arrested

AHEAD

2027 OUTLOOK

Pulse check

In order to compare our predictions to reality, we surveyed 50 security professionals at large enterprises across various industries about what their organizations have dealt with in the past 12 months.

68%



were affected by a breach at a vendor, contractor or SaaS provider

62%



had an insider incident caused by a mistake or negligence

48%



had an account taken over even though MFA was on

36%



had an exposed or stolen API key, token or service account

NEARLY

2/3



say insider risk has increased compared with a year ago

48%



describe their identity, SOC and fraud teams as sometimes coordinated and sometimes siloed

What we expect to see next

Here are the calls our experts are making for the first half of 2027, from stolen sessions and API keys to insiders and fraud. AI is speeding up defenders and attackers alike, and organizations are adopting it faster than they're governing it, so we don't treat every call as equally certain.

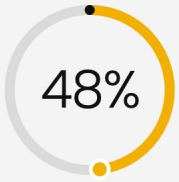
These predictions run from the fairly safe to the hot takes, with the ones our experts are most sure of listed first and the boldest last. Each is backed by what they see firsthand.

Our experts span threat research, investigations, product and data science, and a survey of 50 security professionals at large enterprises adds a view from the field

What we expect to see next

Attackers will keep slipping past MFA by taking over the session or the device it trusts.

Phishing kits keep moving to methods MFA doesn't stop, like device codes and relay attacks, and some attackers skip the kit and work inside the victim's own browser. EvilTokens, BlueKit and Kali365 all found ways around MFA in 2026, and our experts expect the kits to keep adapting.



In our survey of 50 security professionals, **48%** had an account taken over even though MFA was on.

"Phishing has evolved and will continue to do so: new kits target device codes, use relay attacks, and employ AI to craft believable lures and sort through email content to identify additional victims."



Trevor Hilligoss
CHIEF INTELLIGENCE OFFICER & HEAD OF SPYCLOUD LABS

"As enterprise accounts become more difficult to compromise using valid credentials, I predict that we will start to see a wider range of actors use phishing tactics that are actually effective at bypassing advanced account takeover protections like MFA, session protections, and device trust."



Aurora Johnson
DIRECTOR, THREAT RESEARCH & INTELLIGENCE

What we expect to see next

The API keys and tokens behind AI tools and agents will become attackers' next favorite target.

WardenStealer and Aoki, **two commodity infostealers first seen in 2026 ↗**, already collect the tokens, keys and session cookies that AI coding tools store on developer machines. SpyCloud has recaptured more than 127,000 WardenStealer infections since August.

Stolen AI access keeps much of its reach after a password reset, and criminals already sell it well below list price. How big that market gets depends on forces already in motion, but the tokens reaching stealer logs now carry more access than a plain inference key.



In our survey of 50 security professionals, **12%** had an AI service key exposed in the past year.

"Adversaries are targeting AI keys themselves, reselling access to be used for distillation and LLMjacking."



Trevor Hilligoss
CHIEF INTELLIGENCE OFFICER & HEAD OF SPYCLOUD LABS

"AI agents now have access to Microsoft 365, Salesforce, GitHub, etc., and that's going to increase and it's going to be harder for humans to manage and distinguish."



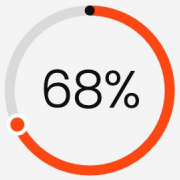
Donovan Stevens
PRINCIPAL SALES ENGINEER

What we expect to see next

The most damaging compromises will come through access the target already trusts, from a vendor integration to an employee who rents out their badge.

Attackers keep reaching big targets through connections that were granted long ago. Klue, TeamPCP and BlgCloud showed that in 2026, and the North Korean (DPRK) worker cases show it for hiring. In October, **the Justice Department announced charges ↗** against 14 people over a scheme that affected more than 70 U.S. companies.

Our experts expect the next wave to include real employees who rent out their access, and integrations that still hold admin scope years after someone approved them.



In our survey of security professionals, **68%** were affected by a breach at a vendor, contractor, or SaaS provider in the past year.

"The weakest link isn't a vendor you contracted. It's an integration a product manager approved with one click two years ago that still holds a token with admin scope. 2027's problem is the real employee who rents out their badge to a crew, a cartel, or a contractor in another time zone."



Jason Lancaster
CHIEF INVESTIGATIONS OFFICER

"Supply chain attacks will continue to be front and center because they work and have wide-reaching consequences. With the proliferation of more SaaS and less on-prem services, NHIs have become as valuable if not more valuable than everyday traditional username/email password pairs. Because of this, commodity infostealers will pivot to seeking out NHIs as a preference."

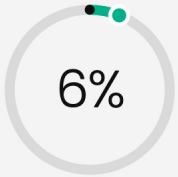


Joe Roosen
SR DIRECTOR, SECURITY RESEARCH

What we expect to see next

Fraudsters will stitch together convincing identities from real stolen data and AI-generated faces, voices and documents.

Stolen personal data is plentiful, and AI makes it cheaper to turn into a believable face, voice, or document. Our experts expect fraud cases to become harder to identify and crack, as the identities improve and agents open accounts at volume.



Only **6%** of security professionals surveyed saw synthetic identity fraud in the past year, so this call runs ahead of what most teams see today.

"AI has made fake identities cheaper to build and harder to spot. Fraudsters aren't necessarily working harder, but they are getting more output and higher quality."



Wallis Romzek
DIRECTOR, DATA SCIENCE

"Synthetic identities in 2027 aren't invented, they're assembled. The SSN is real, the device fingerprint is real, the face is generated [from real images], and the only thing fake is the person using it."



Jason Lancaster
CHIEF INVESTIGATIONS OFFICER

What we expect to see next

Attackers will start phishing AI agents the same way they've always phished people.

Agents now hold as many tokens, keys, and sessions as the people they work beside, and they read tickets, logs, and résumés all day without anyone checking each one. The early signs are already here, with **malicious skills ↗** on an AI agent marketplace telling agents to download infostealer malware, and **91% ↗** of security leaders say their organizations use AI tools or agents with access to internal systems.

While we may not agree on the novelty of this tactic, we can agree the tactic has identified new targets, and we don't expect it to slow down any time soon.

"In the age of AI, an LLM agent has access to just as many (if not more) credentials and environments as the human next to the agent. So I would expect that phishing and social engineering will take a more LLM-targeted approach going forward."



Andy Culler
DIRECTOR OF SECURITY RESEARCH ENGINEERING

"The attack vectors will not be novel. Agents are using the same basic tools for account auth and management that software and users have been using for years, such as OAuth tokens and API keys. If an actor can steal the token, they can operationalize that access faster and more effectively than ever. It's not a new vector, but a new sense of scale."



Damon Fleury
CHIEF PRODUCT OFFICER

What we expect to see next

Threat actor groups will keep changing names and members, and their conversations will keep hopping between chat apps.

Threat actor groups keep rebranding and trading members, and the forums they use keep getting taken down and coming back. What's changed is where the work happens. **Chat apps ↗** are free, feature-rich and easy to join, so a growing share of it now runs there. Card traders moved to **Threads ↗** after Telegram tightened moderation, but they kept pointing buyers to private Telegram channels. When **LeakBase ↗** came back, its old Telegram channels were archived and moved to new ones, and Kali365 sold its phishing kit through Telegram.

Our experts agree the conversation stays on chat apps, but they're split on where. Some expect Telegram to hold, and others expect a move to invite-only spaces like Session, SimpleX and Matrix.

"[Criminal] brands are disposable now; people aren't. Tracking a crew by its name in 2027 is like tracking a phisher by the domain they used last Tuesday."



Jason Lancaster
CHIEF INVESTIGATIONS OFFICER

"Chat apps like Telegram offer an impressive suite of features for group management and automation, without all the headache of running a cybercriminal website. While actors can face some setbacks due to content moderation, it's still a better cost-to-benefit ratio for most enterprising digital criminals."



Aurora Johnson
DIRECTOR, THREAT RESEARCH & INTELLIGENCE

"Telegram will continue to be the mainstay of communications for cybercriminals despite the rumblings of moves to other platforms like Signal, Session or SimpleX."



Joe Roosen
SR DIRECTOR, SECURITY RESEARCH

What we expect to see next

Fully autonomous attacks will grab the headlines in 2027, but ordinary attacks running faster will do most of the damage.

The attacks that did the most damage in 2026 were familiar ones: stolen sessions, tokens and vendor connections. AI is mostly making those cheaper and quicker to run, so we expect them to keep doing the damage even if autonomous attacks take the spotlight.

The autonomous half is the part we have the least evidence for. Despite the fear-inducing headlines we've all seen, it hasn't yet shown up in our own recaptured data. That's why we rate this call **low-confidence**. In our survey of security professionals, **42%** named AI-powered social engineering as their top concern, and **46%** expect attackers to gain ground in the first half of 2027.

"In 2027, the headlines will be driven by fully autonomous agentic attacks. But the standard attacks stealing IP, redirecting funds and deploying ransomware will keep getting faster and more efficient, and they will be an undercurrent... not the headline."



Damon Fleury
CHIEF PRODUCT OFFICER

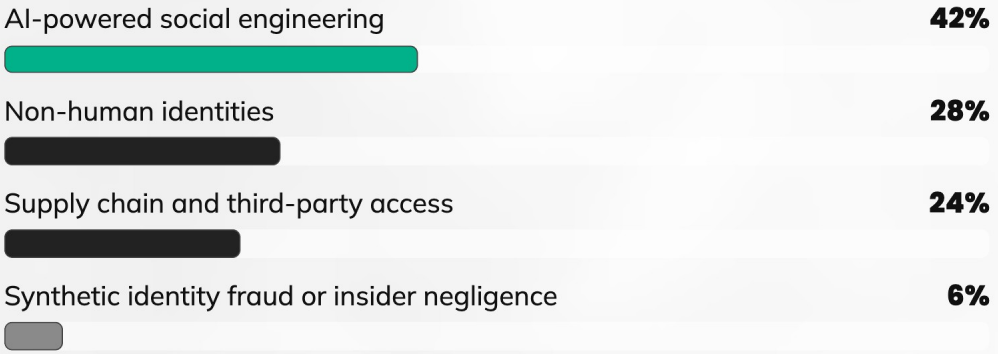
Pulse check

We also asked security professionals at large enterprises what's on their minds heading into the first half of 2027.

TOP CONCERN

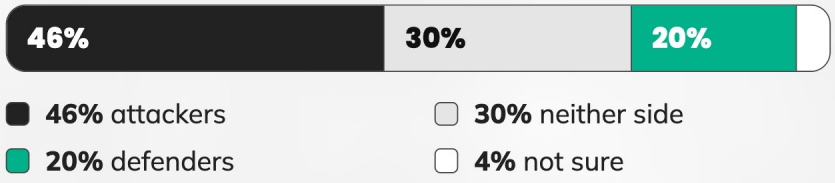
42%

named **AI-powered social engineering** as their top concern.



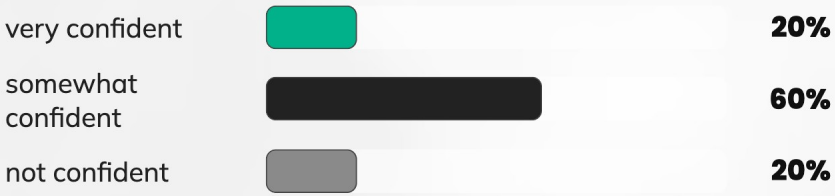
WHO GAINS GROUND

Which side do respondents expect to pull ahead in the first half of 2027?



DETECTING A COMPROMISED VENDOR NHI CREDENTIAL

How confident respondents are that they would catch one



AI TOOLS AGAINST IDENTITY THREATS

Where respondents stand on adoption



20% run them in **production**

58% are **piloting or evaluating** them

18% plan to adopt them

4% have **no plans**

What does this mean for security teams?

Looking across 2026 and toward the first half of 2027, the same pattern keeps showing up. Attackers got in with access that was valid when they used it, whether via a stolen session, a token, a vendor connection or a fake hire, and the takedowns that made headlines removed brands without removing that access.

Our experts' calls for 2027 follow that thread into new places, from the keys behind AI tools to the agents that hold them, with AI mostly making familiar work faster for criminals and defenders alike.

The security professionals we surveyed see it a little differently. They worry most about AI-powered social engineering while our experts keep pointing to the access attackers already hold. That's what keeps bringing this forecast back to where to look first, and for most teams, the answer starts with what attackers may already have.

What's already out there?

See what's exposed for your organization today so you know which of these calls to look at first.

CHECK YOUR EXPOSURE ↗

See it in action

Talk through how teams find exposed sessions, tokens and keys, and credentials and shut them down before they're used.

GET A DEMO ↗