
CISO GUIDE 2027

CISO Guide 2027: Understand Your Social Engineering Risk. Then See How You Rank Against Your Peers.

Your board will not ask whether you ran training in 2027. It will ask how exposed you are, and how you compare. This CISO Guide 2027 shows how to measure your social engineering risk against the threats working in the field, see where you rank against your sector, and align awareness training to the threat model behind the score.

JT

By Jason Thatcher, Founder and CEO, Breacher.ai

**Know your risk. Know where you rank.
Then train for the threats you actually face.**

CISO Guide 2027: the plan in one paragraph

THE SHORT ANSWER

For 2027, start by **understanding your social engineering risk and how you rank against your peers**. Model the threat as the consequential actions an adversary wants, measure how far an orchestrated simulation gets and how many people go with it, and place that score against your sector. Then align security awareness training to your threat model and to your own policy, procedure and process, aimed at the paths your score says are weakest. The outcome to train for is **process invocation**: the employee stops and runs the verification step, not a user who spots the fake. Report your risk band and peer rank to the board, not click rate.

We run orchestrated social engineering simulations against enterprise organizations every week, across email, voice, Microsoft Teams and video, with AI voice agents that hold a live two-way conversation. This guide is the program we would put in front of a CISO building next year's plan. For the threat outlook itself, read our [top 3 emerging AI social engineering risks for 2027](#). For the deepfake-specific playbook, see the [CISO deepfake defense guide](#).

The adversary got faster in 2026. The good news for 2027: the control that beats it is cheap, testable and already yours to deploy.

What changed going into 2027

Three shifts define the threat model a CISO plans against in 2027. None of them is a new technique. Each one removes a constraint the adversary used to live with.

01 / CHANNEL

Voice moved into the tools people trust

Microsoft's 2026 Digital Defense Report puts the weekly volume of confirmed malicious voice phishing over cross-tenant Teams calls up 502% from the same time last year. We break that down in our [Teams vishing analysis](#).

02 / SCALE

AI agents removed the staffing ceiling

A voice campaign used to be limited by how many fluent callers a crew could staff. An agent that holds a live conversation can place and receive calls in parallel, in any language, around the clock.

03 / LIKENESS

Deepfakes learned to talk back

The dangerous deepfake is no longer a pre-recorded clip. It is an interactive likeness on a call, a Teams meeting or a job interview, answering questions in real time. The FBI warned as early as 2022 that deepfakes were appearing in interviews for remote jobs.


The pattern underneath all three is [synthetic normalization](#): plausibility beats fidelity. Microsoft's Q2 2026 data shows technical support impersonation as the dominant lure on Teams, with more than half of Teams phishing attacks in June using generic display names rather than obvious IT support names. The request does not need to look perfect. It only needs to fit a workflow your people already trust, which is exactly why the workflow, not the disguise, belongs at the center of your threat model.

What got worse in 2026: agentic AI made mass voice phishing work

Of the three shifts, one moved further than any other this year. An AI voice agent can now run the entire voice phishing sequence on its own: place the outbound call, leave the voicemail, take the inbound callback, and hold a live two-way conversation that answers questions and steers toward the request. Then it does the same thing with an entire workforce at once.

We know because we run exactly these agents in authorized simulations, and we have measured the change. A year ago the results were mixed. Roughly 1 to 3 percent of users

contacted fell for an AI voice agent, and most people could still tell they were talking to a machine. Across our most recent agentic voice simulations, the weighted mean action rate is 11.5 percent of users contacted, with a range from 0 percent to 34.5 percent of users contacted.



Three things changed, and each one is visible in the field:

Scale

A voice campaign used to be capped by how many fluent callers a crew could staff. An agent places and receives calls in parallel, in any language, around the clock. A small action rate across a whole directory is a large number of people.

Callback

About 11 percent of a contacted user base calls the agent back, and about 34 percent of those who call back take an action. Only about one in nine contacts calls back, but that leg is where the most severe outcomes happen, and the agent answers every one of those calls, at any hour and any volume.

Plausibility

The request is routine, so the red flags awareness training taught (pressure, urgency, suspicion) never appear. People no longer need to believe the voice is human. It only has to sound like support.

Outside data shows the channel growing even before you count agents. In its 2026 Digital Defense Report, Microsoft puts the weekly volume of confirmed malicious voice phishing over cross-tenant Teams calls up 502% from the same time last year, and found that 93% of the attackers it studied succeeded at least once in keeping a target on the line for 20 seconds or longer, the point where a social engineering script can begin. The same report warns that AI can effectively automate social engineering attacks of many types, circumventing language and skill barriers. Mass voice phishing is no longer constrained by labor. That is the single biggest change a 2027 threat model has to absorb, and we break down the full sequence in our [AI voice agent interactive loop](#) findings.

What we are seeing in the field

A threat model should be built from what actually works against organizations, not from headlines. We run these simulations as weekly engagements, and six patterns keep showing up. The full write-up is in our [AI vishing field findings](#).

01

The phone is where people act

Email still reaches more people at first contact. Voice is where they actually divulge information, reset the credential, approve the MFA prompt or open the remote session, because a call reaches a human with no email gateway in between, and on Teams an External label is often the only warning.

02

The callback is the dangerous leg

Most outbound calls go to voicemail. The worst outcomes happen when the employee calls the number back, believing they started a routine interaction. In our experience, most simulations stop before that leg. We break it down in [the callback is the simulation](#).

03

Help desk impersonation is the pretext that works

A call from IT support is routine and expected, so it presents a false green flag. There is no urgency and no pressure for red-flag training to catch, which is why the same pattern runs through the Scattered Spider playbook.

04

Enterprise and outsourced support fare worst

Smaller organizations know their support staff by name. Large enterprises, and anyone with outsourced IT support, cannot tell a real technician from an impersonator by voice, because the technician on the line is usually a stranger.

05

A workforce that invokes process defeats our red teams

Users who spot fakes do not stop us consistently, because suspicion rises and falls with the quality of the pretext. When a target stops a convincing voice or deepfake, they almost never do it by spotting the fake. They stop, verify through a separate channel, and the sequence ends there. We detail this in [verification procedure training](#).

06

The gap is a process gap

Organizations rarely fail for lack of tools or lessons. They fail where no written step exists, or where the check ran on a channel the caller supplied. In our engagement data, how seriously an organization takes security matters more than the platform it runs.

Read those six together and the conclusion is uncomfortable for a catalog training program. The pretexts that work in the field carry no red flags, arrive on channels the training never covered, and are defeated by process rather than perception. A program aligned to that threat model looks very different from one aligned to last year's phishing email, and so does the risk score that measures it.

Why generic awareness training is not the answer in 2027

Security awareness training has real value. It builds a reporting culture, explains the threat, and teaches people to pause on unusual requests. The problem in 2027 is not awareness. It is fit.

Classic red-flag training teaches people to watch for pressure, urgency and an unfamiliar voice. The help desk callback we run in the field presents none of them. The employee dialed the number. The request is routine. Support is outsourced, so an unfamiliar voice is normal. Every cue reads safe, a false green flag, and the training designed to catch the scam never fires.

Your verification step is not anyone else's. Your approvers, your help desk ticketing rules, your vendor onboarding flow and your wire thresholds are specific to you. A catalog course written for every organization cannot teach the callback your accounts payable team is supposed to make, because the vendor that wrote it has never seen your procedure.

Catalog content also plays catch-up: by the time a new pretext becomes a library module, it has already been used against someone.

That is why training has to be **aligned to your policy, procedure and process**, and to the security controls you actually run. Those are three different layers, and generic content touches none of them.

POLICY, PROCEDURE, PROCESS

Training only holds when it teaches all three layers of your own control, using one example: a callback required to a known number.

POLICY

The rule

Every credential reset, payment change or remote session is verified by calling back a known number, never one the requester supplies.

PROCEDURE

The steps

Which number, who calls, what to say, what to do if the requester objects, and where the outcome gets recorded.

PROCESS

The workflow

Built into the help desk ticket, new-hire provisioning and vendor banking changes, so the step cannot be skipped quietly.

A catalog module can teach the idea of a callback. Only training built from your own documents can teach your number, your steps and your workflow.

Hyper-training users is the wrong answer too

The opposite reflex is just as costly. When the threat gets worse, many programs respond with more: monthly modules for everyone, longer courses, and automatic retraining after every click. That volume does not add control. It teaches people to click through training the same way they click through warnings, and it spends attention on topics most roles will never face.

HYPER-TRAINING

More of everything, for everyone

The same catalog content pushed to every employee on a calendar, sized to the vendor's library rather than your risk. Fatigue rises, attention falls, and no specific control gets stronger.

TAILORED TRAINING

The right procedure, for the right role

Short modules generated from your own policy, procedure and process, sent to the people who perform each consequential action, aimed at the paths your risk score shows are weakest.

The goal is not more training. It is less training that maps to more control.

SPOTTING THE FAKE, OR INVOKING THE PROCESS

Both try to stop a synthetic request before a consequential action. Only one keeps its value as generation quality improves.

PERCEPTION

Spotting the deepfake

The target has to recognize the voice, face or message as fake. Its value is set by the adversary's generation quality, which improves every quarter.

Decays on the adversary's schedule

PROCEDURE

Invoking the process

A callback to a known-good number or a second-channel confirmation never evaluates the fake at all. It asks the same question of a crude fake and a perfect one.

Holds regardless of fidelity

Illustrative trajectories, not measured values. The full argument is in [deepfake detection training is a decaying control](#).

No bank relies on a teller's eye alone to catch a counterfeit note. It requires the check. The bill got better. The check did not have to. That is why the best deepfake security awareness training in 2027 teaches the verification step and its trigger, not a list of visual artifacts the next model release will remove.

THE DESIRED OUTCOME

Process invocation is the moment an employee stops before a consequential action and runs the verification step, whatever they believe about the request. It does not depend on noticing anything wrong. That is the outcome every module, simulation and score in this guide is built around, and it is the one behavior that still works when the fake is perfect.

**A workforce that invokes process defeats our red teams.
Users who spot fakes do not, at least not consistently.**

This is an argument about where human judgment sits, not against technology. Machine-level identity verification belongs in the threat model, especially in hiring and customer onboarding: document verification, liveness checks, validation against authoritative sources and synthetic media screening on the interview channel, with a hard gate before any account is provisioned. Those controls deserve investment, and they deserve testing to prove they hold.

Precise about the claim. A better fake does not break a verification step directly. It raises the pressure to wave the step through. That is why a 2027 program measures the exception rate, the share of required checks that are skipped or satisfied on the wrong channel, rather than claiming any procedure is unbreakable.

Understand your risk, and how you rank against your peers

A CISO cannot prioritize what cannot be measured, and cannot defend a budget without context. Two numbers do both jobs: your own social engineering risk, and where it sits against organizations like yours.

YOUR RISK

The OSES™ Score

A [social engineering risk score](#) built from two terms reported side by side. Depth asks how far the simulation got: the deepest point any single person reached. Spread asks how many people went with it. Together they land on a four-band scale: Low, Medium, High or Critical.

YOUR RANK

Your sector benchmark


The [Social Engineering Risk Index](#) computes sector medians from real engagement outcomes, not survey answers, and every engagement places your score against your sector's median, so the board sees whether you are ahead of your peers or behind them.

The two terms matter because a single number gets risk backwards. An organization where almost nobody engaged, but one person handed credentials to an external endpoint, looks excellent on a click report. An organization where many people moved but nothing followed looks like a crisis. Depth catches the first, Spread catches the second, and the band is never reported without both. We publish [how the score is built](#) and [how the index is computed](#), because a benchmark a CISO cannot defend in front of a board is not worth having.

Why peers matter. “We scored Medium” starts a debate. “We scored Medium and our sector median is High” ends one, in either direction. Rank turns a risk number into a decision about where the next dollar goes.

How to align awareness training to your threat model

The program runs as a loop. You model the threat, measure your exposure to it, see where you stand against your sector, and point training at what the score found.



The threat model is not a list of scary scenarios. It is a list of **consequential actions**: an outbound payment, a vendor banking-detail change, a credential reset, an MFA enrollment, a privileged access grant, a software install on request, a data export, a new account provisioned for a hire. For each one, ask which channels can reach the person who performs it, which pretext an adversary would use, and what step should fire before they act.

Threat model to training map

Here is how six 2027 threats translate into a procedure to train and a way to test it. Use it as a starting template for your own.

THREAT MODEL TO TRAINING MAP

Each row pairs a 2027 threat with the action it targets, the procedure to train, and how to test that the procedure holds.

2027 THREAT	ACTION IT TARGETS	PROCEDURE TO TRAIN	HOW TO TEST IT
Teams help desk vishing	MFA reset, remote session	Open a ticket and call back on the published service desk number. Never start a session from an unsolicited chat.	Microsoft Teams impersonation simulation
Deepfake executive on a call	Outbound payment	Call back on the number already on file. Dual approval above the payment threshold.	CEO fraud simulation
Vendor banking detail change fraud	Vendor bank account change	Confirm with the vendor through the number in the vendor master, never the one in the request.	Deepfake phishing simulation
Voicemail and inbound callback, run by AI voice agents	Credential reset, remote session, data disclosure	Reach support only on the static, documented service desk number. Never call back a number left in a voicemail.	AI vishing simulation , callback included
Deepfake job candidate	New account provisioned	Machine identity verification clears before IT provisions anything. Recruiters cannot waive it.	Fraudulent hire scenario
Synthetic identity at onboarding	Customer account opened	Document and liveness checks with a manual escalation path that cannot be talked around.	Deepfake KYC penetration testing

Rows are examples. Your own map comes from your consequential actions and the channels that reach the people who perform them.

Notice what the map does to training once your score shows which rows are weakest. A finance clerk and a service desk analyst no longer receive the same module. Each receives the procedure they are actually asked to perform, in the channel they are actually reached on. That is what risk based security training should mean: training sized to the risks your threat model ranks highest. It is also what makes behavior based security training measurable, because each module points at a behavior a simulation can observe.

We built the tool to make that happen

Aligning training to a threat model used to mean months of custom content work. That is the gap we built the [Breacher.ai Secure Behavior Management platform](#) to close. It generates awareness training and matching simulations tailored to your organization and aligned to your security controls, in minutes, aimed at the paths your risk score says are weakest. Fewer modules, each one mapped to a control.

FROM YOUR POLICY

Drop in a procedure

Upload a wire approval procedure or a help desk verification standard. The module comes back in your terminology, naming your systems, your approvers and your callback steps.

FROM WHAT FAILED

Point it at a finding

When a simulation shows a step did not hold, Studio generates the remediation module for that exact step and routes it to the roles that perform it.

FROM AN ARTICLE

Drop in the threat

Read about a new pretext this morning? Drop the article in and the platform generates both the simulation and the matching training, the same day the threat moves.

FROM SCRATCH

Describe the objective

Tell the AI Concierge what people should be able to do after the module. It drafts the content and the knowledge check for your team to shape in Studio.

Training generates in any language, every element stays editable, and modules export as SCORM. The [AI phishing simulation generator](#) builds the matching scenario from the same

source, so what people learn and what they are tested on line up exactly. It is the [AI-powered social engineering training platform](#) we wished existed when we started running engagements, and we explain the design choices behind it in [how we built an AI-powered awareness training platform](#).

It fits the stack you already run: SOC 2, SSO, Entra ID and Google, plus an API, a CLI and webhooks. Already invested in Microsoft? The [add-on for Microsoft Attack Simulation Training](#) syncs with Entra ID and layers AI voice and deepfake video on top of what you run today. It is currently in beta, starting with a free 90-day pilot.

Bring one procedure from your threat model. Watch it become a module, a knowledge check and a matching simulation.

[Book Your Demo](#)

The 2027 CISO agenda: seven moves

Each move builds on the one before it. The first two need no simulation at all, so a CISO can start on day one of the new budget year.

01

Write the threat model as consequential actions

List the actions that move money, access or data: outbound payments, vendor banking changes, credential resets, MFA enrollment, privileged access grants, software installs on request, data exports and new-hire provisioning. Then map the channels and pretexts that reach the people who perform each one.

02

Measure verification coverage before anything runs

For each consequential action, ask whether a defined verification step exists, whether it is channel-independent, and whether the person asked to act can reach it under time pressure. A verification coverage assessment needs no simulation and shows where the uncovered paths are.

03

Bind verification to the request, not the requester

The check fires on the request class, whoever appears to be asking: a callback to a known-good number, confirmation through the real ticket portal, or a second approver. A channel the requester supplied never counts, because every channel the adversary hands over is one the adversary controls.

04

Measure your risk under the pressure the field shows

Run an orchestrated social engineering simulation that chains email, SMS, voice, Teams and video, where each stage adapts to what the target did before it, with the voicemail and inbound callback included. Score it on Depth and Spread so you know how far it got and how many went with it.

05

Rank against your peers

Place your risk band against your sector median from the Social Engineering Risk Index. Rank tells you whether your exposure is normal for organizations like yours or a problem the board needs to fund.

06

Train for what you find

Generate training from your own procedures for the paths your score says are weakest, and route it to the roles that perform them. Key every remediation module to the control that did not hold. The module a person receives is the procedure they did not follow.

07

Report your band and rank, not click rate

Take the board your risk band, your peer rank and the Depth and Spread behind them, with the coverage gaps and the training plan that follow. That is a conversation about exposure and investment, not about who clicked.

Where the 2027 budget goes

Use this as a sanity check on next year's line items. The question for each is simple: does it keep its value as the fakes improve?

2027 BUDGET CHECK

Six common social engineering line items, rated on whether they hold as generation quality improves.

LINE ITEM	WHAT IT BUYS IN 2027	AS FAKES IMPROVE
Procedural verification coverage	A channel-independent check on every consequential action, owned by the process, not the person.	● Holds
Orchestrated simulation, live voice and video	Evidence that each check fires under realistic, multi-stage pressure, including the inbound callback.	● Holds, scores the rest
Training generated from your threat model	Modules on the exact steps your people perform, in minutes, in any language.	● Holds
Machine identity verification	Document, liveness and authoritative-source checks with a hard gate before provisioning.	● Holds, test it
General security awareness training	Reporting culture, threat context and the habit of pausing on unusual requests.	● Holds its role
Training people to spot deepfakes	A perceptual skill the adversary's next model release is designed to beat.	● Decays

Green: keeps its value as generation quality improves. Amber: value declines from the day it is delivered.

What a CISO should report to the board in 2027

Boards asked about deepfakes in 2026. In 2027 they will ask whether the organization would survive one. Click rate cannot answer that, and in voice, Teams and video channels

there is no link to click. Lead with your band and your peer rank, then show the process invocation rate that explains them. For framing the conversation itself, see [how CISOs can answer the board about deepfakes](#).

RISK BAND**How exposed are we?**

Your OSES™ Score on the four-band scale, Low to Critical. The number that goes in front of the board.

PEER RANK**How do we compare?**

Your band against your sector median from the Social Engineering Risk Index. Context the board can act on.

DEPTH AND SPREAD**How far did it get, and how many went?**

The two terms behind the band. Narrow and deep is a different problem from broad and shallow, and the board should see both.

PROCESS INVOCATION RATE**Did people run the step?**

Verification executed through an independent channel, over verification required, also reported as process hold rate. The outcome training exists to produce.

COVERAGE RATE**Is every consequential action covered?**

Actions with a defined verification step, over all consequential actions identified. The number most organizations have never calculated.

EXCEPTION RATE**How often was the step waived?**

Checks consciously skipped or rerouted, over checks required. The variable that moves as the fakes improve.

Per-user results still matter. They route training to the people and process paths that need it. They are a routing input, not a board metric, because a score on a person does not tell you whether the wire goes out. Human risk management measures who is likely to fail. This program measures whether the organization survives when they do, which is why it sits comfortably above an HRM tool you may already own.

A 90-day plan for Q1 2027

Every step below starts with assets you already have: your process documents, your help desk, your finance team and your identity stack.

DAYS 1 TO 30

Model the threat

- List every consequential action across finance, IT, HR and identity.
- Map the channels and pretexts that reach the people who perform them.
- Run a verification coverage assessment to see which paths have no written step.

DAYS 31 TO 60

Score and rank

- Run an orchestrated simulation across voice, Teams and email, callback included.
- Get your OSES™ Score, with Depth and Spread reported side by side.
- Place your band against your sector median.

DAYS 61 TO 90

Train and report

- Generate training on the procedures behind your weakest paths.
- Write verification steps for the uncovered paths, starting with the help desk.
- Take the board your band, your peer rank and the plan that follows from both.

By the end of the first quarter, the CISO walks into the board meeting with a threat model, a risk band, a peer rank and a training plan aimed at the weakest paths. That is a fundamentally stronger position than a click-rate trend line.

How Breacher.ai runs the 2027 program

Breacher.ai was built for exactly this program. The OSES™ engine runs authorized, fully automated, orchestrated simulations across email, SMS, voice, Microsoft Teams and video. Its AI voice agents hold a live two-way conversation and handle the inbound callback, and orchestration means stage two changes based on what the target did at stage one, the way real adversaries operate. It is a sequence most simulation tools were not built to run end to end. Every finding is named against the procedure that failed, never against a person.

The [simulation platform](#) covers the full threat model in one program: [deepfake phishing simulation software](#) for voice, video calls and email, [voice phishing simulation](#), [help desk impersonation simulation](#), [CEO fraud prevention](#) testing and [executive tabletop simulation](#). Need a scoped, point-in-time social engineering penetration test instead of a program? Our [deepfake red team engagement](#) is scoped to your workflows and reported against the procedures that failed.

Your people are not the weak point in 2027. They are the ones executing the procedure. Give them a procedure built for the threats they actually face, and give your board a risk number it can compare.

Measure your risk. Then train for it.

Thirty minutes. We will map your consequential actions, show how your OSES™ Score is built, and where your sector median sits today.

[Book Your Demo](#)

Frequently asked questions

What should be in a CISO guide for 2027?

A CISO guide for 2027 should cover what is working against organizations in the field, how to measure social engineering risk, how to benchmark it against peers, how to align awareness training to the threat model, where to put budget and what to report to the board. The core is understanding your risk and where you rank, then training for process invocation against the threats you actually face.

How do you measure social engineering risk?

Run an orchestrated simulation that ends at named consequential actions and score it on two terms. Depth is the deepest point any single person reached, so one consequential action is enough to raise it. Spread is the share of the population that moved. The Breacher.ai OSES Score reports both side by side and places the result on a four-band scale: Low, Medium, High or Critical.

How do you benchmark social engineering risk against your peers?

Compare your risk band with the median for your sector. The Breacher.ai Social Engineering Risk Index computes sector medians from real engagement outcomes rather than survey answers, which makes it a social engineering benchmark a CISO can use to tell the board whether the organization is ahead of or behind comparable organizations, and why.

How do you align security awareness training to your threat model?

Start from the consequential actions an adversary wants, such as payments, vendor banking changes, credential resets and new-hire provisioning. Map the channels and pretexts that reach the people who perform each one, write the verification step that should fire, and generate training on that exact step for those roles, built from your own policy, procedure and process. Let your risk score decide which paths get trained first.

What are you seeing in the field going into 2027?

Voice is where people take the consequential action, the inbound callback is where the worst outcomes happen, and help desk impersonation is the most effective pretext because it presents no red flags. Large enterprises and organizations with outsourced IT support fare worst. People who defeat a convincing voice or deepfake almost always do it by invoking process, not by spotting the fake.

How effective has AI voice phishing become in 2026?

In Breacher.ai engagement data, a year ago roughly 1 to 3 percent of users contacted fell for an AI voice agent. Across our most recent agentic voice simulations, the weighted mean action rate is 11.5 percent of users contacted, with a range from 0 to 34.5 percent of users contacted. About 11 percent of a contacted user base calls the agent back, and about 34 percent of those who call back take an action. Because an agent needs no human caller, it runs that sequence against an entire workforce at once.

Should CISOs still fund deepfake awareness training in 2027?

Yes, when it trains the verification procedure. Training people to spot deepfakes by eye and ear is a decaying control, because its value falls every time generation quality improves. Deepfake security awareness training that teaches the callback, the ticket check and the second approver keeps its value, especially when it is generated from your own procedures and aimed at the paths your risk score flags.

What outcome should security awareness training aim for?

Process invocation: the employee stops before a consequential action and runs the verification step, whatever they believe about the request. It is the desired outcome instead of a user who spots the fake. In our red team engagements, a workforce that invokes process defeats the simulation, while users who spot fakes do not do so consistently.

Is more security awareness training the answer?

No. Generic training is not the answer, and hyper-training users is the wrong answer too. Pushing more catalog modules to everyone on a calendar builds fatigue without strengthening any specific control. Training works when it is tailored to your organization, aligned to your own policy, procedure and process, and sent only to the roles that perform each consequential action.

What is the difference between a social engineering simulation and a social engineering penetration test?

A social engineering simulation measures behavior across a population and produces a risk score you can benchmark against peers. A social engineering penetration test is a scoped, point-in-time engagement that proves whether a specific control can be bypassed. Breacher.ai runs both, as simulation programs and as scoped red team engagements, and reports each finding against the procedure that failed rather than the person involved.

How do you prevent CEO fraud in 2027?

Require a callback on the number already on file and a second approver above a payment threshold for every outbound payment and vendor banking change, whoever appears to be asking. Then test it with a deepfake CEO fraud simulation that uses a cloned voice on a live call, because a procedure that has never faced realistic pressure is an assumption.

What social engineering metrics should a CISO report to the board in 2027?

Lead with your risk band and your peer rank, then show the terms behind them: Depth, Spread, process invocation rate, coverage rate and exception rate. Together they describe how exposed the organization is, how it compares with its sector and where the gaps are. Click rate does not, and voice, Teams and video channels have no link to click.

How does Breacher.ai generate training aligned to our threat model?

The Breacher.ai Secure Behavior Management platform generates awareness training and matching simulations in minutes from four starting points: a policy or procedure, a simulation finding, an article about a new threat, or an objective described to the AI Concierge. Training generates in any language, stays editable in Studio and exports as SCORM, aimed at the paths your risk score says are weakest.

Sources

- [Microsoft Digital Defense Report 2026 \(report page; vishing figures in the full report, July 2025 to June 2026\)](#)
- [Microsoft Digital Defense Report 2026, full report \(PDF\)](#)
- [Microsoft Security Blog: Email threat landscape, Q2 2026 trends and insights, July 23, 2026](#)
- [CISA Advisory AA23-320A: Scattered Spider, help desk social engineering](#)
- [FBI IC3 PSA I-062822-PSA: Deepfakes and Stolen PII Utilized to Apply for Remote Work Positions, June 28, 2022](#)

Jason Thatcher



Founder and CEO of Breacher.ai and creator of OSES™. Fifteen years in security operations and offensive testing, previously at ZeroFox, Deepwatch, and GuidePoint Security. He builds and runs orchestrated social engineering simulations against enterprise organizations.

Find Out Where You Rank Before Your Board Asks

Bring the consequential action that worries you most, a wire, a vendor banking change, an MFA reset or a new hire. We will show you how we would score your exposure to it, where your sector lands, and the training that follows from the result.

[Book Your Demo](#)