FOR CISOS AND SECURITY LEADERS

Know Your Risk. Map It. Close the Gaps.

Ransomware crews, deepfake impersonators and AI voice agents all come through the same door: a person who can reset a password, move money or grant access. Inside: the field data, the OSES™ Risk Score, a risk grid, a People, Process and Technology plan, and training generated from your own procedures.

WHAT THE FIELD DATA SHOWS

AI voice phishing now works at scale

2025

1 to 3%

of users contacted fell for AI voice phishing

VS

2026

11.5%

average action rate, of users contacted, across our most recent simulations

THREE FINDINGS THAT SHOULD CHANGE YOUR 2027 PLAN

FINDING 01

The callback is the danger zone

Most outbound calls go to voicemail. The inbound callback is where the most severe outcomes happen, because the user called **us**, so there is no urgency and no suspicion. A voice test that does not capture the callback misses the most important result.

FINDING 02

IT support is a false green flag

A password reset call from support is routine. Support is often outsourced, so an unfamiliar voice is normal. Red-flag training never fires because nothing looks wrong. People read identifiers aloud on a call that they would never type into a form.

FINDING 03

Process beats perception

Users who try to spot the fake do not stop our red teams consistently. Users who invoke process do. One sentence ends the simulation: **"I need to verify this request with our internal team."**

THE TOP 3 DEEPFAKE HEISTS OF 2026

$400M+

FaceStealer, North Korea

Crypto firms and developers · disrupted by the DOJ, Nov 2026

Recorded webcam footage and deepfakes of trusted contacts on fake Zoom and Teams calls talked victims into installing malware. More than $400M stolen over 18 months.

Process that stops it: verify any meeting invite out of band, and never install software because a caller asks.

$129M

The AI CEO and lawyer ring

Thousands of companies targeted · SDNY indictment, Nov 2026

An e-signature NDA lure, then a call with an AI clone of the CEO voice and at times face, then a fake lawyer from a major firm instructing the wire. Alleged, July 2024 to February 2026.

Process that stops it: call the CEO back on the number on record, and require dual approval on any confidential deal.

€95M

Fideuram, Intesa Sanpaolo

Private bank, Italy · February 2026

A WhatsApp message posing as the group CEO, then a call with an AI clone of a lawyer the chairman knew. Eleven transfers abroad. About €36M was still missing when it became public.

Process that stops it: verify through a channel the requester did not choose, no matter how familiar the voice.

Sources: US Department of Justice, Nov 12 2026 (FaceStealer); US Attorney SDNY indictment unsealed Nov 2026, charges are allegations and defendants are presumed innocent; Reuters and Il Sole 24 Ore, Sep 2026 (Fideuram, about $108M). Every one of the three ended at a consequential action that a channel-independent verification step would have stopped.

UNDERSTAND YOUR RISK, AND HOW YOU RANK AGAINST YOUR PEERS

The OSES™ Score: measure the control, not the click

A click rate tells you who was curious. It does not tell you whether the wire went out. The OSES™ Score measures how far an orchestrated simulation got and how much of your workforce went with it, then places you on a four-band scale you can report to the board.

OSES™ Score = Depth + Spread

DEPTH

How far did it get?

The deepest point any single person reached. Not divided by headcount, because one person is enough to be breached. Depth floors at 70 the moment an outcome produces a real consequence, so a breached organization can never score low.

SPREAD

How many went with it?

The share of the population that moved. Per capita and capped, so volume can never mask a contained breach, and a click counts for a fiftieth of a consequential action.

THE ESCALATION LADDER

L0	No response to the lure at all.
L1	Opened the email, answered the call, heard the voicemail. Scores zero.
L2	Clicked. Reached the page or payload and went no further.
L3	Engaged. Replied, conversed or called back, without acting.
L4	Took a consequential action: credentials, payment, access, data.
L5	Escalation. The action produced a real consequence.

Each person is counted once, at the deepest rung they reached.

FOUR BANDS

Low	Medium	High	Critical
0 to 39	40 to 59	60 to 79	80 to 100

Band boundaries sit at 40, 60 and 80, so the Depth floor of 70 always lands an escalation in High or above.

HOW THE TERMS BEHAVE: PUBLISHED EXAMPLES

SECTOR	SCORE	BAND
Technology	72.7	High
Manufacturing	69.8	High
Financial services	56.8	Medium
Legal	46.7	Medium

DEPARTMENT	SCORE	BAND
HR	88.4	Critical
All-staff	83.0	Critical
Finance	80.4	Critical
IT	66.5	High

Source: breacher.ai/oses-risk-index/. Published to show how the two terms behave, not to state where any one sector sits. Department figures are a separate cut of the same data and do not sum with the sector table.

WHAT TO REPORT TO THE BOARD

Coverage rate
Consequential actions with a channel-independent verification step, over all consequential actions you identified.

Process hold rate
Simulated requests where verification was executed, over simulated requests that required it.

Exception rate
Simulated requests where verification was waived or bypassed, over requests that required it. Watch it as fakes improve.

MAP YOUR RISK

The consequential action risk grid

Map risk on the actions an adversary wants, not on the lures they send. Place each consequential action by how likely it is to be targeted and how much damage one success would cause. The top-right cells are where verification has to hold first.

	LOW	MEDIUM	HIGH LIKELIHOOD
HIGH IMPACT	LOW LIKELIHOOD Privileged access grant	MEDIUM LIKELIHOOD Outbound payment or wire Vendor banking-detail change Fake candidate hire	HIGH LIKELIHOOD · ACT FIRST Help desk password reset MFA re-enrollment or phone change Remote support session
MEDIUM IMPACT	LOW LIKELIHOOD Physical access grant	MEDIUM LIKELIHOOD Data export or disclosure	HIGH LIKELIHOOD Identifier read-back on a call
LOW IMPACT	LOW LIKELIHOOD	MEDIUM LIKELIHOOD	HIGH LIKELIHOOD Link click, no action

Placement is the Breacher.ai default assessment for a typical enterprise, not field data. Move each item to fit your environment.

SCORE YOUR COVERAGE TODAY

CONSEQUENTIAL ACTION	TYPICAL PRETEXT AND CHANNEL	RISK	VERIFICATION STEP DEFINED?	OWNER
Help desk password reset	Caller poses as an employee or executive, often with a cloned voice	Critical	☐ Yes ☐ No	
MFA re-enrollment or phone change	"Lost my phone before a board call"	Critical	☐ Yes ☐ No	
Remote support session	Fake IT support over Teams or phone	Critical	☐ Yes ☐ No	
Outbound payment or wire	Executive voice or video on a call, e-signature lure first	High	☐ Yes ☐ No	
Vendor banking-detail change	Spoofed vendor email followed by a confirming call	High	☐ Yes ☐ No	
Fake candidate hire	Deepfake interview, then onboarding and provisioning	High	☐ Yes ☐ No	
Privileged access grant	Urgent request from a senior engineer or contractor	High	☐ Yes ☐ No	
Identifier read-back	"Read me your badge or employee number to confirm"	Medium	☐ Yes ☐ No	
Data export or disclosure	Auditor, regulator or partner pretext	Medium	☐ Yes ☐ No	

Your coverage rate is the number of Yes answers over the number of rows. Most organizations cover outbound wires and little else.

PEOPLE, PROCESS AND TECHNOLOGY

How to close the gaps on every action

Every consequential action needs all three layers. Technology reduces how often a person has to decide. Process makes the decision channel-independent. People rehearse the step until invoking it is a reflex.

ACTION	PEOPLE	PROCESS	TECHNOLOGY
Help desk password reset	Agents rehearse refusing resets that skip verification, even for executives	Callback to the number on record, never the number offered	Self-service reset with document and liveness checks as the default path
MFA re-enrollment or phone change	Agents treat every MFA change as a privileged action	Second-channel approval from the manager or account owner	Phishing-resistant MFA (FIDO2, passkeys) so lost-factor cases are rare
Remote support session	Employees know support never cold-calls for remote access	Verify any support contact through the service portal	Restrict external Teams contact and unapproved remote tools
Outbound payment or wire	Finance rehearses the callback under executive pressure	Dual approval plus callback for any new or urgent payment	Payment platform holds for new payees and changed amounts
Vendor banking-detail change	AP staff know the change request is the lure	Confirm with a known vendor contact from the master file	Bank account validation and change-hold window in the ERP
Fake candidate hire	Recruiters and hiring managers know the pretext is a candidate	Identity gate before offer and again before provisioning	Document verification, liveness and synthetic media detection on the interview channel
Privileged access grant	Admins escalate, never fulfil, urgent access requests	High-assurance or in-person verification for Tier 0 accounts	Just-in-time access with approval workflow and logging
Identifier read-back	Everyone knows no one legitimate asks for an ID aloud	Never confirm identifiers on an inbound or unsolicited call	Identifiers alone never authenticate a request
Data export or disclosure	Data owners verify the requester, not the story	Written request through a known channel plus owner sign-off	DLP and approval gates on bulk export

HOW BREACHER.AI CLOSES THE LOOP

SIMULATE

Orchestrated, like real adversaries

OSES™ sequences email, voice, Teams and live deepfake video with a context layer: an email, a wait, then a call that references it. AI voice agents handle the outbound call, the voicemail and the inbound callback. Multi-channel is not orchestration.

SCORE

Findings against procedures

Coverage, process hold and exception rates by action, plus your OSES™ Score and where you rank against your peers. Findings name the control that held or failed, not the person.

TRAIN

Built from your own policies

Secure Behavior Management generates training from your policies, procedures and simulation results in minutes, in any language, aimed at the paths your score says are weakest.

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

AWARENESS TRAINING, GENERATED FROM YOUR PROCEDURE AND PROCESS

Train the step you wrote, not the fake you cannot see

Generic awareness training cannot teach a callback rule it does not know exists. Breacher.ai Secure Behavior Management reads your own policies, procedures and simulation findings and generates the training that rehearses them, in minutes, in any language. Every module teaches one thing: invoke the process.

01 · INPUT

Your procedure

Upload the help desk reset SOP, the payment approval policy, your acceptable use policy, or drop in a threat report. Or start from your simulation findings.

→

02 · GENERATE

Training in minutes

The platform turns your exact steps into micro-modules, scenarios and matching simulations. The AI Concierge turns a news story into both a simulation and the training for it.

→

03 · TARGET

The paths that failed

Modules go to the roles and procedures your OSES™ Score says are weakest: the service desk, finance, HR. No hyper-training everyone on everything.

→

04 · MEASURE

Process hold rate

The next orchestrated simulation shows whether the step now holds under pressure. Coverage, process hold and exception rates go to the board.

EXAMPLE: WHAT ONE GENERATED MODULE LOOKS LIKE

GENERATED FROM: IT-SOP-014 HELP DESK CREDENTIAL RESET

The Callback Rule: Verifying Every Reset Request

Audience: Service desk agents · Length: 6 minutes · Assigned after: inbound reset simulation

1. Your rule, in your words

"No password or MFA reset without a callback to the number in the HR directory. Never use a number the caller provides." Quoted from your SOP, not a generic template.

2. The scenario

A caller who sounds exactly like your CFO says she is locked out minutes before a board call and asks you to read out a temporary password. The voice is perfect. The procedure does not change.

3. The rehearsal

Agents practise the exact line: "I will call you back on the number we have on file." Then they practise holding it when the caller pushes back, escalates or gets angry.

4. The check

Three questions on when the callback applies, which number to use, and what to do when it fails. Pass rate and completion route into the platform report.

Illustrative module. Content is generated from the customer's own documents; the SOP name is an example.

GENERIC TRAINING

Teaches the tell

Watch for lip sync, odd cadence, urgency. Decays as fakes improve, and never mentions your callback rule, your approval chain or your escalation path.

GENERATED FROM YOUR PROCESS

Teaches the step

Your rule, your systems, your people, rehearsed until invoking it is a reflex. Works the same against a crude fake and a perfect one, and can be measured.

Also available: import your own content, sync users with Entra ID or Google, run alongside Microsoft Attack Simulation Training, SSO and API access. SOC 2.

Find out where you rank before your board asks.

Run an orchestrated simulation, get your OSES™ Score and a gap map for every consequential action.

Book Your Demo → breacher.ai/service-demo