PLAYBOOK

AI Voice Phishing

Running a voice phishing programme that reaches every employee, measures verification and trains on the call itself.

Callstrike's AI caller phones your employees in their own language, listens and answers, and records who shared and who refused. This playbook sets out how to plan a campaign, walks through three programmes from start to finish, and shows how awareness leads, partners and compliance owners each run it.

01 READ FIRST

Consent, limits and disclosure

Every campaign is awareness training for your own employees. These commitments hold for every call.

Consent first

A campaign calls only employees with consent on record. Consent Management is in every plan, so no programme has to go without it.

Stop when asked

The caller ends the call when an employee asks it to stop, hang up or end the call, or says the workspace's stop word.

Lines never crossed

Every scenario is checked against a content policy before it is saved, and threats, extortion and content built to distress are refused.

Tell every employee

The wizard has a step for how employees learn it was a simulation: in-call education, a follow-up email, or both. Choose at least one for every campaign.

A VOICEMAIL NEVER HEARS THE PRETEXT When a call reaches voicemail, the caller leaves nothing, or one short fixed message you choose. Calls answered by a phone's screening assistant are marked screened and never scored as a conversation.

Five decisions, in the order the wizard asks them

1 **The scenario**

Start from the Official library, where difficulty is measured on five dimensions and shown as Obvious, Realistic or Convincing, or draft a custom scenario from three questions.

2 **After the call**

Decide how employees learn it was a simulation, and which follow-up each outcome receives. Employees who shared are enrolled in remediation training.

3 **The audience**

Start with a department or a country, not the whole company. A first campaign is a measurement, so keep it small enough to read every transcript.

4 **Call settings**

Spread the calls over weeks inside a daily window, choose the pace, and check the caller number for each audience country. Recording can be turned off, and call length capped.

5 **Test, then launch**

Play the employee in a browser test call. The caller waits for you to speak first, as it will on a real call.

> **RULE OF THUMB** Pick the difficulty for the question you are asking. Obvious tells you who will act on anything. Convincing tells you whether your verification routine holds when the call sounds right.

Three programmes, from start to finish

Each example names who runs it, who is called and what a strong response looks like. The pretexts are the everyday requests employees already receive.

EXAMPLE 1 · RUN BY THE AWARENESS LEAD

A first baseline: the service desk asks for a sign-in code

All departments · Three countries · Difficulty: Realistic

THE SETTING	A manufacturer with offices in the UK, Germany and France wants a first measurement before rolling out new verification guidance.
THE CALLS	Each country runs as its own campaign in its language, English, German or French. Calls come from a local number in the employee's country, spread over three weeks of working hours.
STRONG RESPONSE	The employee declines to read out a code, ends the call and reports it to the service desk on the number they already have.
AFTERWARDS	Employees who shared are enrolled in remediation training and the others receive the follow-up you chose. Employees not reached go into a second round you launch.

EXAMPLE 2 · RUN BY AN MSSP FOR A CLIENT

Payroll asks to confirm bank details before the month-end run

Finance and payroll | Client workspace | Difficulty: Convincing

THE SETTING	A managed security provider runs quarterly voice tests for a retail client and reports to the client's head of security.
THE CALLS	The partner launches from the client's workspace, where calls are limited to the client's verified domains while the partner manages it.
STRONG RESPONSE	The employee refuses to change or confirm bank details on an inbound call and raises it through the payroll portal.
AFTERWARDS	The partner reports results by department, with each call's believability score and transcript as evidence.

EXAMPLE 3 · RUN BY THE AWARENESS LEAD, REPORTED TO THE RISK COMMITTEE

The same scenario a quarter later, one level harder

Repeat measurement | Customer service | Difficulty: Obvious to Convincing

THE SETTING	After training, the team repeats the courier delivery scenario that caught most of customer service, at a harder difficulty.
THE CALLS	The same scenario is re-run from the library, and the dashboard tracks each employee's risk score across both campaigns.
STRONG RESPONSE	The employee asks for a reference and offers to call the courier back on its published number.
AFTERWARDS	Report the trend by department to the risk committee, never a list of names, and pick the next scenario from the weakest result.

04 FOR AWARENESS AND HUMAN RISK LEADS

Measure the habit, then train on the call

Voice phishing tests the moment an employee decides whether to trust a caller. Measure that moment and train on it.

- Run a baseline at Realistic difficulty first, then move one level at a time.
- Score the behaviour you trained: did the employee pause, decline and call back on a number they already had?
- Use each employee's own transcript in their training: it shows what gave the call away.

- Treat no answer, voicemail and screened calls as not reached, never as a pass, and give them a second round.

> **WHAT TO WATCH FOR** A department whose results improve only at the easiest level. The habit has not formed yet.

05 FOR CONSULTANCIES AND MSSPS

Run it for every client from one account

Partners manage each client's workspace from one partner account, under their own branding.

1 **Agree scope**

Agree the departments, countries and difficulty with the client's security lead in writing.

2 **Check consent and domains**

Confirm consent on record for every employee in scope. While you manage the workspace, calls are limited to the client's verified domains.

3 **Launch and follow up**

Run the campaign in the client's working hours and keep the disclosure step switched on.

4 **Report**

Report by department with the believability scores and transcripts behind them.

> **WHAT TO WATCH FOR** A client who asks to call suppliers or customers. They never consented, so they are never called.

06 FOR COMPLIANCE AND AUDIT

Evidence that the programme runs as written

Auditors ask whether a control is tested, how often, and what happened next. A voice programme can answer all three.

EU notice

An EU AI Act notice turns on automatically for employees in the EU.

Audit log

A tamper-evident log of admin actions, exported to CSV.

Consent

Only employees with consent on record are called.

> **WHAT TO WATCH FOR** A programme that tests once a year. Agree a cadence, quarterly for high-risk roles, and keep each campaign's results with the policy it tests.

A checklist for launch day

Sign-off

The programme owner, legal and HR have agreed the scenario, the audience and the window.

Service desk briefed

The service desk knows a campaign is running and how to answer an employee who reports a call.

Test call done

You have played the employee in a browser test call and heard the scenario end to end.

Disclosure chosen

In-call education, the follow-up email, or both, are switched on for this campaign.

See it on a call with your own team. Book a walkthrough at www.callstrike.ai

AI VOICE PHISHING · NOVEMBER 2026