Stopping Help-Desk Impersonation: How to Close the Attack Surface That Brought Down MGM and Caesars

People & In-Person / Last updated 2026-06-11 / https://www.scrambleid.com/learn/stopping-helpdesk-impersonation-with-people-verification

> **Status (June 2026):** Early access. The People verification family is live with early-access customers and isn't generally available yet. This article describes the shipped design as it stands today; talk to your ScrambleID account team about access and timelines before gating help-desk procedure on it.

In one sentence: Help-desk impersonation has driven some of the largest breaches of the past three years; knowledge-based questions and callback-to-known-good are not adequate defenses in the AI era; person-to-person cryptographic verification turns the help desk from the highest-leverage attack surface in the enterprise into a deterministic verification gate that AI cannot defeat.

TL;DR (canonical)

- **Why this matters:** The help desk has the authority to bypass every other authentication control (reset passwords, re-enroll MFA, add devices, grant privileged access). Compromising the help desk is fast, high-impact, and historically under-defended.
- **Why traditional defenses no longer work:** Knowledge-based questions (KBA) are derivable from breached data and LinkedIn. Callback-to-known-good is defeated by SIM swap and voice cloning. Manager confirmation is defeated by deepfake voice on the manager call. These were defenses for a pre-AI threat model.
- **What works:** Person-to-person cryptographic verification at every help-desk action that materially changes account state (password reset, MFA re-enrollment, device add, privileged access). The verification takes a few seconds end-to-end (versus the 30 to 90 seconds typical of knowledge-based questions), is deterministic, and is AI-quality-independent.
- **The policy that makes this work:** No sensitive help-desk action proceeds without a successful people verification, full stop. The procedural exception, "I'll just verify you with these questions because people verification isn't working," is how the entire defense fails.
- **The recovery path matters as much as the primary:** A weak recovery flow becomes the new attack surface. Cold recovery requires identity proofing plus dual control, not a help-desk

override.

Why the help desk is the highest-leverage attack target

Three years of breach disclosures have established the pattern. Scattered Spider compromised MGM Resorts and Caesars Entertainment in 2023 by impersonating employees to the IT service desk and walking the agents through credential resets and MFA re-enrollment. The 2025 Scattered Spider wave ran the same playbook against UK retail, insurance, and aviation targets. LAPSUS$ ran variations across multiple targets. The technique works because:

1. **The help desk has by-design authority.** Resetting a password, adding an MFA device, or granting privileged access is not an exception in help-desk operations. It is the daily work. The help desk cannot operate by refusing every state-changing request.
2. **Attackers prepare meticulously.** OSINT from LinkedIn (manager names, recent projects, employee IDs from social posts), prior breach data (employee email, phone, address, SSN, prior passwords), and public corporate communications give attackers what they need to pass KBA-style verification.
3. **Voice cloning is now commodity.** The "callback to a known number" defense assumes the voice on the callback is a meaningful signal. Voice cloning from 30 seconds of recorded audio is now production capability. The voice on the callback is no longer evidence of legitimacy.
4. **Time pressure works.** Attackers introduce urgency ("I'm in a customer meeting and locked out, can you please reset this now?") that pushes help-desk agents past the procedural friction that's supposed to stop them.
5. **Help-desk turnover is high.** Tier-1 service desk roles see frequent staff change. Training decays. New agents inherit habits. Procedural rigor varies day to day.
6. **The blast radius is enormous.** Once the attacker is logged in as an executive or a privileged engineer, they have everything that account has. Lateral movement is hours, not days.

The current generation of help-desk attacks is fast (typically 15-45 minutes from first call to compromise), high-impact (often credential or admin access), and increasingly hard to detect because the attacker is using legitimate-looking interactions.

What attackers actually do

The typical help-desk impersonation attack follows a pattern. Understanding it is the first step in defending against it.

Stage 1, target selection. The attacker picks an employee with valuable access (executive, finance leader, IT engineer, sales operations). They harvest LinkedIn for the employee's manager, role, recent projects, location, tenure. Prior-breach data fills in employee ID, phone, prior addresses.

Stage 2, pretext development. The attacker builds a plausible reason for the call. "I'm traveling and my MFA token isn't working." "I just got a new phone and need to re-enroll." "I'm locked out and have a board meeting in 20 minutes." Time pressure is added to discourage the help-desk agent from procedural caution.

Stage 3, the call. The attacker calls the help-desk number from a spoofed caller ID matching the employee's known number, or uses a number-spoofing service to display the corporate caller ID. Voice cloning of the executive (if available) is layered in for the more sophisticated operators. The attacker walks the agent through a credential reset, MFA re-enrollment, or device add.

Stage 4, immediate exploitation. Once the attacker has new credentials and a new MFA device, they immediately log in, deploy persistence, and begin lateral movement. By the time the legitimate employee notices anything is wrong, the attacker has hours to days of dwell time.

Where traditional defenses fail

Each link in the traditional defense chain has a known failure mode in the current threat environment:

Defense	Why it fails today
Security questions (employee ID, manager name, recent project)	All derivable from LinkedIn and prior-breach data
Caller ID match	Spoofable in seconds with consumer tools
Callback to known-good number	Defeated by SIM swap; defeated by call forwarding compromise; defeated by voice cloning
Manager confirmation call	Manager voice can also be cloned; manager may be in a meeting and confirm hastily
Photo on file vs camera selfie	Defeated by AI face augmentation; not always operationally feasible
"I recognize their voice from prior calls"	Voice cloning from 30 seconds of audio defeats this
Procedural delay (24-hour cooldown)	Attackers know the cooldown and target before-cooldown moments
Self-service password reset over email	Becomes the attack surface (compromised email = compromised account)

None of these are useless. They remain weak signals worth retaining as background context. The point is that no combination of these alone is sufficient verification for an action as consequential as credential reset or MFA re-enrollment for a high-value account.

The people verification gate

The architectural change is to require a deterministic cryptographic verification at every help-desk action that materially changes account state. The flow:

1. Employee contacts the help desk (call, chat, ticket).
2. Help-desk agent confirms the request and identifies what action is being requested (password reset, MFA re-enrollment, device add, privileged access grant).
3. Agent initiates a people verification request through the employee's enrolled identity.
4. Employee's mobile authenticator prompts. Employee approves on their phone (Face ID, Touch ID, or device PIN as required by policy).
5. Verification completes in a few seconds. Both parties see the result. The agent sees the employee's verified identity card with the appropriate work attributes.
6. Agent proceeds with the requested action. Audit log captures the people verification event with the device IDs (ZIDs), the attribute set, and the timestamp.

If the verification fails or times out (60-second TTL), the agent does not proceed. The attacker has no private key to sign with; the verification cannot complete; the action does not happen.

The only way an attacker defeats this is to also possess the legitimate employee's enrolled mobile device, which is a fundamentally different (and much harder) threat than impersonating a phone call.

Action-by-action policy

Different help-desk actions warrant different verification rigor. A useful starting taxonomy:

Action	Required verification	Notes
Read-only inquiry ("what's my account status?")	Soft signals OK (knowledge questions, caller ID)	Low blast radius if compromised
Software install request	people verification	Could be a malware install request
Password reset	people verification + manager attestation for executives	Not just a productivity action; account-takeover-equivalent
MFA re-enrollment	people verification + dual control for privileged accounts	The post-passkey ATO surface
Device add (laptop, phone)	people verification + 24-hour delay for high-risk devices	Persistence vector
Privileged access elevation	people verification + dual control (Lockstep, in development)	Cannot be self-authorized
Account unlock after lockout	people verification	Was the lockout caused by attempted compromise?
Wire-transfer authorization (finance)	People verification + dual control + named-payee confirmation	See Lockstep article
Vendor banking change	people verification of the vendor contact + AP dual control	See Recovery and Fallback Playbook

Action	Required verification	Notes
Personnel record change	people verification + HR concurrence	Insider misuse vector

Each row is a policy choice. The pattern: the higher the blast radius if the request is fraudulent, the more layered the verification. People verification is the floor for any action with material blast radius.

What "people verification gate" looks like in practice

A worked example: an employee calls the help desk to re-enroll MFA after replacing their phone.

- Employee: "Hi, I just got a new phone and need to set up my MFA again. I have an exec briefing in 30 minutes."
- Agent: "Sure, before I do that I need to complete a people verification with you on your enrolled device. Are you with your prior phone, or do you need to use the cold-recovery process?"
- Employee (legitimate): "I have my old phone here, I haven't deactivated it yet."
- Agent: "Great. I'm sending the people verification request now, please open the ScrambleID app and approve."
- (Two seconds later) Verification completes. Employee identity card visible to agent.
- Agent: "Verified. I'll initiate the new-device enrollment now. You'll get a setup link on the new phone."

Or, in the attacker case:

- Attacker: "Hi, I just got a new phone and need to set up my MFA again. I have a board meeting in 20 minutes."
- Agent: "Sure, before I do that I need to complete a people verification with you on your enrolled device. Are you with your prior phone?"
- Attacker: "Uh, my old phone broke. I don't have it anymore."
- Agent: "OK, in that case we need to go through the cold-recovery process, which requires identity proofing plus your manager's attestation. That'll take about 15 minutes. Can you stay on the line, or would you like me to call you back at the number on file?"

The cold-recovery path is intentionally slower and more rigorous. An attacker on a 20-minute schedule will often hang up rather than complete it. A legitimate employee with a genuinely lost phone completes it without issue.

The recovery problem

The most common failure mode is not the primary verification; it is the recovery path that gets used when the primary fails. "Help, I lost my phone, I can't do people verification" turns into "OK, just answer these security questions instead" and the entire architecture collapses.

The mitigation is to design recovery deliberately:

- **Warm path** (employee has a secondary device or another enrolled authenticator): use the secondary directly. people verification still happens, just from the secondary device.
- **Cold path** (employee has no working authenticator): identity proofing process. Government-ID document verification, biometric matching, manager attestation, and a cooling-off period appropriate to the role. See Recovery and Fallback Playbook.
- **Assisted path** (in-office employee, IT can verify in person): verified physical handoff plus manager concurrence. Still produces a signed audit event.

The principle: recovery must be at the same assurance level as the primary. A help-desk-driven password reset that bypasses people verification because the user "lost their phone" is the same attack surface as a help desk that doesn't require people verification at all.

Operational considerations

Help-desk training. Agents need to understand why the policy exists and how to handle the social-engineering pressure techniques (urgency, authority, sympathy). The training that worked in 2018 ("trust your gut") fails against AI-quality voice cloning and prepared OSINT. The training that works in 2026 is procedural: "I cannot complete this action without a successful people verification, regardless of the urgency claimed." There is no judgment call for the agent to make.

Executive carve-outs. Executives often request carve-outs ("this is taking too long, just override it"). These requests are exactly the threat model. Document a clear policy: no executive carve-outs for sensitive actions, including for the executive themselves. Brief executives in advance so they know to expect people verification on every help-desk interaction.

Vendor and contractor coverage. Help-desk impersonation also targets the vendor and contractor channels. Vendors calling Accounts Payable to change banking details. Contractors arriving at facilities. The same people verification primitive applies, with the verification anchored to the vendor or contractor's enrolled identity. See People Verification for Finance: Wire Transfers and Vendor Banking Changes (next in this tier).

SIEM and SOC integration. people verification events are audit records that feed the SIEM. Failed verifications, repeated attempts, and unusual patterns (geographic, time-of-day, action-type) feed the SOC. The combination of people verification at the help desk and behavioral analysis in the SOC is more powerful than either alone.

Tabletop the failure modes. Run incident-response tabletops where the scenario is "Scattered Spider just impersonated our finance leader to the help desk, they got a password reset, MFA re-enrollment, and now have access to the wire-payment system. What broke?" The tabletop will surface the procedural exceptions, the cold-recovery weaknesses, and the executive carve-outs you have not yet closed.

Standards and compliance posture

Help-desk procedures gated by people verification map cleanly to:

- **NIST SP 800-63B** authenticator posture for high-assurance contexts (device-bound keys, WebAuthn UV).
- **NIST SP 800-207 (Zero Trust)** continuous verification and least privilege at every authentication event.
- **CISA phishing-resistant MFA guidance** for the workforce-side authenticator.
- **NYDFS Part 500.12** (financial services) MFA requirements, plus Part 500.7 privileged-account controls, plus Part 500.6 audit.
- **PCI DSS v4.0.1** requirement 8.4 MFA on non-console access into the CDE, plus 7.x access control.
- **HIPAA Security Rule** technical safeguards for ePHI access.
- **SOC 2** CC6.1 logical access, CC6.2 authentication, CC7.x monitoring.
- **ISO/IEC 27001:2022** A.5.16 / A.5.17 / A.8.5 (identity, authentication, secure authentication).

For a worked compliance mapping, see Compliance Mapping: NIST and CISA.

What this is not

Not a substitute for the rest of the security stack. people verification at the help desk closes a specific high-leverage attack surface. EDR, MDM, network segmentation, SIEM, threat hunting, and identity governance all remain necessary.

Not a replacement for managed device controls. A jailbroken or rooted phone fails the device-posture checks that the ScrambleID mobile app enforces, but managed-device hygiene (MDM, EDR on the phone, patching) is still the responsibility of the broader IT program.

Not effective if the help-desk policy has exceptions. The defense is the policy plus the technology, not the technology alone. An organization with people verification deployed and a "but for executives we make exceptions" carve-out has the same exposure they had before people verification.

Not a one-off project. The procedural rigor needs to be maintained. New agents need training. Executive expectations need refreshing. Recovery flows need periodic exercise (drill the cold path quarterly).

Key Takeaway

Help-desk impersonation has driven some of the largest breaches of recent years (MGM, Caesars, and the 2025 Scattered Spider wave across UK retail, insurance, and aviation). The traditional defenses (knowledge-based questions, callback to known-good, manager confirmation, photo on file) were designed before AI-driven social engineering at scale and before commodity voice cloning,

and they are no longer sufficient as primary verification for credential resets, MFA re-enrollment, device adds, or privileged access grants. Person-to-person cryptographic verification turns the help desk into a deterministic verification gate: the employee's hardware-bound private key signs a server-issued, single-use challenge with a 60-second TTL, the verification completes in a few seconds end-to-end (versus the 30 to 90 seconds typical of knowledge-based questions), and the help-desk action proceeds only on success. The attacker who does not hold the legitimate employee's private key cannot complete the round trip regardless of how convincing the voice or context. The policy that makes this work is "no exceptions for sensitive actions," combined with deliberate cold-recovery design (identity proofing plus dual control) so the recovery path does not become the new attack surface.

FAQ

Why is the IT help desk a major attack target?

The help desk has authority to reset passwords, re-enroll MFA, add devices, and provision access. An attacker who successfully impersonates a legitimate user to the help desk gets to bypass every other authentication control. The Scattered Spider group used this pattern to compromise MGM and Caesars Entertainment in 2023; ALPHV/BlackCat, LAPSUS$, and others have used variants. Help-desk social engineering is fast (minutes), high-impact (full account takeover), and historically under-defended.

Aren't security questions and callback verification enough?

No. Knowledge-based questions (employee ID, manager name, recent project, last password change) are derivable from breached LinkedIn data, social media, and prior breaches. Callback to the user's known number is defeated by SIM swap and voice cloning. Both methods were designed before AI-driven social engineering at scale and before deepfake voice was commodity capability. They remain useful as low-assurance fallback signals; they are not sufficient as primary verification for credential resets, MFA re-enrollment, or device adds.

How does people verification work for help-desk requests?

When an employee contacts the help desk for a sensitive action (password reset, MFA re-enrollment, device add, privileged access request), the help-desk agent initiates a people verification through the employee's enrolled identity. The employee's device authenticator (with hardware-bound private key) signs a challenge bound to a single-use Dynamic Identifier with a 60-second TTL. Both parties see the verification complete in real time. The action proceeds only on success. The attacker, who does not hold the legitimate employee's private key, cannot complete the round trip regardless of how convincing the voice or context.

What if the employee genuinely lost their phone?

This is the recovery problem and it must be addressed deliberately or it becomes the new attack surface. The recommended pattern is identity-proofing-based new-device enrollment for cold recovery (see the Recovery and Fallback Playbook), with dual-control approval for high-risk recoveries (Lockstep, in development, is designed to enforce this). The recovery path must be at the same assurance level as the primary; help-desk-driven password resets without identity proofing are how passkey-protected accounts get taken over.

Does this slow down the help desk?

people verification completes in a few seconds once the employee triggers it on their phone. Compared to the 30 to 90 seconds typically spent on knowledge-based question-and-answer, people verification is faster, not slower. The friction is one-time enrollment of the authenticator; once enrolled, every help-desk verification is that fast.

What if the employee's phone is the thing being compromised?

People verification uses hardware-protected key storage (Apple Secure Enclave, Android StrongBox, Windows TPM). On iOS and Android, the private key cannot be extracted by jailbreak tooling. The mobile app blocks cryptographic operations if the device fails posture checks. A compromised phone is a different threat (covered by device posture and conditional access) and not the same as a help-desk impersonation attempt.

Can we apply this to vendor and contractor verification too?

Yes. The same people verification primitive is used for vendor banking changes (Accounts Payable verifying the legitimate vendor before changing payment details), contractor verification at physical sites (security verifying the contractor's enrolled identity), and executive sign-offs (finance verifying an executive request before wiring funds). Help-desk impersonation is one of the highest-leverage applications of people verification, but the primitive applies wherever a human-to-human verification needs deterministic cryptographic proof.

References (public)

- CISA Phishing-Resistant MFA Implementation Guidance: https://www.cisa.gov/sites/default/files/publications/fact-sheet-implementing-phishing-resistant-mfa-508c.pdf
- FBI / CISA Joint Advisory on Scattered Spider (AA23-320A): https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-320a
- FinCEN Alert on Deepfake Media for Identity Fraud: https://www.fincen.gov/sites/default/files/2024-11/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf

- NIST SP 800-63B: https://csrc.nist.gov/pubs/sp/800/63b/4/final
- NIST SP 800-207 (Zero Trust): https://csrc.nist.gov/publications/detail/sp/800-207/final

---

Related reading

- What Is People Verification?
- People Verification vs Photo ID, Video, Notary, and KBA
- Deepfake-Resistant Identity Verification
- Recovery and Fallback Playbook
- Lockstep: Dual Control
- People Verification for Finance: Wire Transfers and Vendor Banking Changes
- Compliance Mapping: NIST and CISA