

Deepfake Detection and Prevention: A Cybersecurity Guide

Practical Methods to Identify and Prevent AI-Generated Impersonation

A technical reference for cybersecurity professionals, IT teams, fraud prevention teams, and organisational security stakeholders.

GSDC CYBERSECURITY SERIES

ENTERPRISE EDITION

Understanding Deepfake Attacks

Deepfakes are AI-generated media — audio, images, or video — engineered to convincingly replicate the appearance or voice of a real person. Unlike crude video editing of the past, modern deepfakes are produced using generative AI models capable of synthesising highly realistic facial movements, voice cadence, and emotional expression. The result is synthetic content that can, in many cases, pass initial human scrutiny.

In a cybersecurity context, deepfakes are not an abstract risk. They are an active attack vector. Threat actors use them to impersonate executives in fraudulent payment requests, to clone voices for social engineering over phone calls, and to fabricate video evidence in identity-based fraud. The attack surface extends to any interaction where identity is established through audio or visual cues — which includes a significant portion of modern business communication.

The challenge for organisations is that traditional identity verification methods were not designed to account for AI-generated impersonation. Caller ID, visual recognition in a video call, and even voice familiarity are all exploitable signals when an adversary has access to sufficient source media. A single publicly available conference recording or social media video can provide enough raw material to generate a convincing synthetic version of an individual's face and voice.

Why Traditional Controls Fall Short

- Visual recognition in video calls can be spoofed with real-time face-swap tools
- Voice familiarity is compromised by voice cloning from public recordings
- Caller ID and email sender addresses are easily spoofed independently
- Human cognitive bias towards authority figures increases susceptibility

The Threat Landscape

Deepfake-enabled attacks combine social engineering with synthetic media to bypass procedural controls. The threat is especially acute for finance teams, helpdesks, and anyone authorising high-value actions based on verbal or visual instructions from senior personnel.

⚠ Any process that relies on recognising a voice or face as the sole verification method is vulnerable to deepfake exploitation.

Common Types of Deepfake Attacks

Deepfake attacks are not monolithic. They manifest across multiple communication channels and exploit different organisational processes. Understanding each attack type — and what a realistic scenario looks like — is essential for building targeted defences.

Executive Impersonation & Payment Fraud

A synthetic video or audio message from a "CEO" instructs a finance employee to execute an urgent wire transfer to a new account, citing a confidential acquisition or regulatory deadline. The request bypasses normal approval channels because it appears to come directly from leadership.

Example: An employee receives a WhatsApp voice note from what sounds like the CFO, requesting an immediate transfer of £80,000 to a supplier. The voice is cloned from public earnings call recordings.

Voice Cloning & Fake Audio Instructions

Attackers clone a known individual's voice using readily available audio samples and use it to issue instructions over phone or messaging platforms. The cloned voice is indistinguishable from the real person to untrained ears.

Example: A branch manager receives a call appearing to be from their regional director, instructing them to override a security hold on a customer account. The director's voice was cloned from publicly posted video interviews.

Video Conference Impersonation

Using real-time deepfake tools, an attacker replaces their face with that of a known executive or partner during a live video call. The impersonation is used to authorise actions, extract sensitive information, or build false trust.

Example: During a Zoom call, a "senior partner" requests that a junior lawyer share confidential client documents. The partner's face and voice are synthesised using real-time AI tooling.

Helpdesk & Account Recovery Fraud

Synthetic audio or a fabricated ID document is used to impersonate an employee or customer when contacting a helpdesk or account recovery process. The goal is to reset credentials or gain access to privileged systems.

Example: A threat actor calls the IT helpdesk using a cloned employee voice, claims to be locked out of their account, and requests a password reset without completing secondary verification.

Vendor & Customer Identity Impersonation

Deepfakes are used to impersonate trusted vendors or customers in procurement, contracting, or customer service contexts — enabling fraud, data theft, or manipulation of business relationships.

Example: A procurement officer receives a video message from an apparent long-term supplier requesting a change of bank account details before a large invoice is due. The video is a deepfake generated from the supplier's LinkedIn profile photo and public videos.

How Deepfakes Are Created

Understanding the technical mechanisms behind deepfake creation helps security teams appreciate both the sophistication of the threat and its practical limitations. Deepfakes are produced using generative AI — a class of machine learning models designed to create new, synthetic content that resembles training data.

Generative AI Architectures

Generative Adversarial Networks (GANs) consist of two competing neural networks — a generator that creates synthetic media and a discriminator that attempts to identify it as fake. Through iterative training, the generator learns to produce increasingly convincing output. GANs are widely used for facial synthesis and image manipulation.

Diffusion models represent a newer approach, generating high-quality images by progressively refining random noise into structured content. They are increasingly used for photorealistic face and scene generation.

Transformer-based audio models are used for voice cloning and speech synthesis. Given a few seconds to a few minutes of a target's voice, these models can generate new speech in that voice — including arbitrary words and phrases not present in the original recording.

Creation Techniques

- **Facial manipulation:** Swapping or blending a target's face onto a source video body using encoder-decoder networks trained on paired facial data
- **Lip-sync generation:** Modifying mouth movements in existing video to match new audio, enabling fabricated speech to be placed in authentic-looking footage
- **Voice cloning:** Training a voice model on target audio samples to synthesise new utterances in the target's voice and style
- **Real-time face swap:** Applying face replacement in live video streams using lightweight models optimised for low-latency inference

From Source Media to Synthetic Content

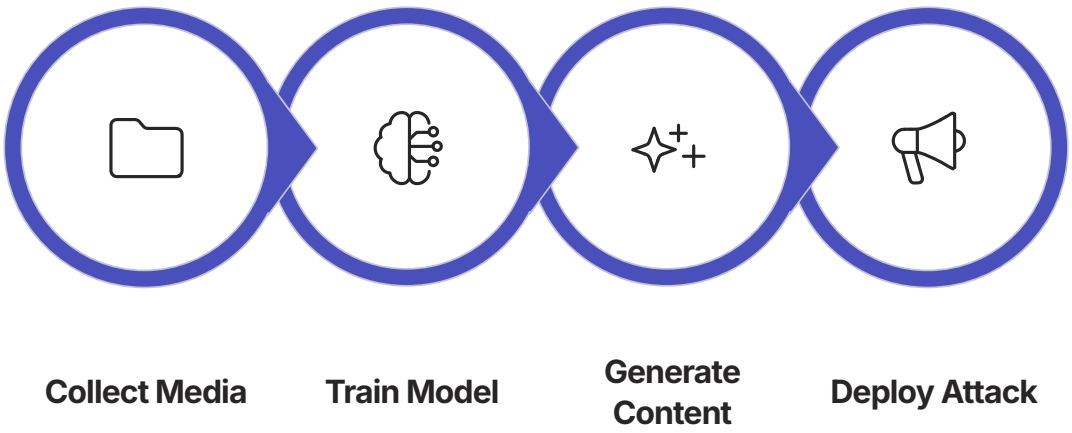

⚠ Publicly available media — LinkedIn profiles, conference recordings, earnings calls, social media videos — provides sufficient source material for voice cloning and facial deepfake generation without requiring any direct access to the target.

Identifying Visual Deepfakes

Visual deepfake detection requires systematic analysis of specific facial and environmental signals. While AI-generated video quality continues to improve, current generation tools still produce characteristic artefacts that trained observers and automated systems can identify. No single indicator is conclusive on its own — detection confidence increases when multiple anomalies are present simultaneously.



Irregular Blinking & Eye Movement

Early deepfake models were notorious for producing faces that rarely blinked. While more recent models have addressed this, inconsistent blink frequency, asymmetric eye movements, or pupils that do not track naturally with the subject's head movement remain indicators of synthetic generation.



Facial Boundary Artefacts

The transition zone between a synthesised face and the original background or hairline often exhibits blurring, colour banding, or unnatural edge softness. Examine the jaw, ear, and hairline boundaries carefully — these regions are frequently poorly rendered, especially under non-frontal angles.



Lighting & Shadow Inconsistencies

Deepfake models synthesise facial texture and lighting based on training data, which may not match the environmental lighting in the source video. Watch for highlights on the nose, forehead, or cheeks that are inconsistent with the apparent light source, or shadows that do not fall in the expected direction.



Lip-Sync Mismatches

When a deepfake is generated by modifying a video to match new audio, the lip movements may not perfectly synchronise with the spoken words. Look for slight delays, over- or under-articulation of consonants, or mouth shapes that do not match the phonemes being spoken.



Unnatural Head & Facial Expressions

Realistic human facial expression involves coordinated micro-movements across the full face — forehead, eyebrows, cheeks, and chin. Deepfakes frequently flatten or homogenise these micro-expressions, producing a face that appears emotionally static or whose expressions are limited to the mouth and eye region only.



Skin Texture & Detail Loss

Generated faces often exhibit unnaturally smooth skin with reduced fine detail — reduced pore visibility, absent facial hair at natural density, or anomalous reflectivity. Under high-resolution examination or zoom, synthesised skin lacks the stochastic variation present in real human skin captured on camera.

These visual indicators are most useful during post-incident forensic review or when reviewing recorded media. Real-time detection during a live call requires dedicated software tools rather than manual observation alone.

Detecting AI-Generated Audio

Voice cloning and AI-generated speech present a distinct detection challenge. Unlike visual deepfakes, which can be reviewed frame by frame, audio impersonation often occurs in real-time over phone calls or voice messages — leaving limited time for analysis during the interaction itself. Effective audio deepfake detection combines human awareness with automated signal processing.

Human-Observable Indicators

- **Unnatural pauses:** Cloned speech may insert or omit natural micro-pauses between words, phrases, or at sentence boundaries — particularly when the AI model is generating novel sentences not present in training data
- **Flat prosody:** Emotional colouring — the natural variation in pitch, pace, and emphasis that conveys sentiment — is often reduced or inconsistent in synthesised speech
- **Pronunciation anomalies:** Unusual stress patterns on specific words, or subtle mispronunciations of proper nouns, place names, and industry-specific terminology that may not have been well-represented in training data
- **Background noise inconsistencies:** Cloned audio may lack consistent ambient sound, or the background noise may cut in and out unnaturally — particularly when segments of speech are stitched together from multiple generated clips
- **Unusual verbal habits:** Absence of a known person's characteristic filler words, verbal tics, or distinctive phrasing patterns

Technical Analysis Methods

Waveform & Spectrogram Analysis

Forensic audio tools can reveal discontinuities in waveform energy distribution, unnatural spectral smoothness, or artefacts in frequency bands typically occupied by breath sounds and room noise. AI-generated speech often exhibits anomalously clean spectrograms absent the natural irregularities of real vocal production.

Voice Biometrics & Speaker Verification

Speaker verification systems compare audio samples against enrolled voiceprint templates. These systems calculate probabilistic similarity scores rather than binary matches and can flag audio that does not meet threshold confidence levels for the claimed speaker identity.

Detection Limitations

No audio detection method is fully reliable in real-time. High-quality voice clones trained on substantial audio samples can pass basic perceptual tests. Speaker verification systems can be fooled by clones trained on the enrolled speaker's own samples.

Audio Analysis Workflow

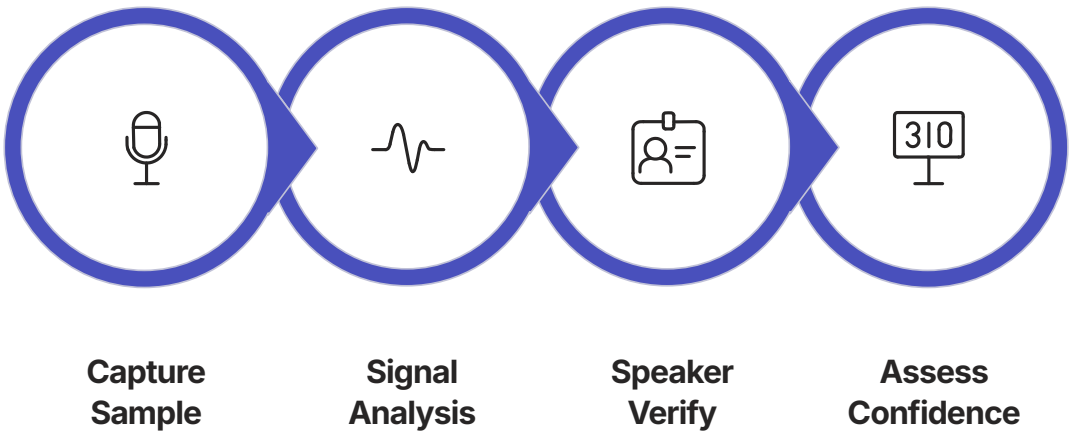

AI-Powered Deepfake Detection Techniques

Manual inspection alone is insufficient for reliable deepfake detection at scale. AI-powered detection tools apply machine learning to identify patterns in media that are imperceptible to human observers. These tools operate by analysing multiple signals simultaneously and producing confidence scores that indicate the likelihood of synthetic generation.

CNN-Based Image & Video Analysis

Convolutional Neural Networks (CNNs) extract spatial features from individual video frames — detecting artefacts in texture, compression patterns, and facial geometry that are characteristic of GAN or diffusion model output. Frame-level analysis is then aggregated across a clip to produce a video-level assessment.

Transformer-Based Detection Models

Vision transformers analyse global image context — capturing long-range relationships between facial regions that CNN models may miss. They are particularly effective at identifying inconsistencies in lighting, pose, and background that only become apparent when considering the full image composition.

Facial Landmark Analysis

Detection models map 68 or more facial landmark points per frame and analyse geometric consistency across time. Deepfakes frequently introduce subtle distortions in landmark positions — particularly around the eyes, nose, and mouth — that deviate from the natural physics of human facial movement.

Temporal Consistency Analysis

Analysing consistency between consecutive video frames reveals flickering, inconsistent textures, or discontinuous facial expressions that are introduced by frame-level synthesis. Models trained on temporal data are more sensitive to these inter-frame artefacts than single-frame classifiers.

Audio-Video Synchronisation

Cross-modal analysis tools examine the synchronisation between audio signals and visual lip movements. Deepfakes that splice new audio onto existing video, or modify lip movements to match new speech, frequently introduce measurable synchronisation offsets detectable by dedicated models.

Combining Detection Signals

No single detection method is definitive. Robust deepfake detection architectures combine outputs from multiple classifiers — visual, temporal, audio, and cross-modal — and aggregate confidence scores using ensemble methods. The resulting composite score, combined with contextual signals such as request urgency and communication channel, informs the final risk assessment.

Media Forensics and Content Provenance

Media forensics and content provenance tools provide an additional layer of analysis beyond AI detection classifiers. Rather than analysing what the media depicts, provenance techniques examine how and where the media was created, modified, and transmitted — looking for signs of tampering or absent authenticity credentials.

Metadata & File Integrity Checks

Digital media files contain embedded metadata — including creation timestamps, software application identifiers, device information, and GPS coordinates where applicable. Forensic examination of metadata can reveal inconsistencies: a video claiming to have been recorded yesterday but containing a creation timestamp from two years ago, or metadata indicating the file was last modified by a video editing application.

File integrity checks compare hash values against known-good references where available, and examine compression signatures for evidence of re-encoding — a common artefact when deepfake-generated media is embedded into existing legitimate footage.

Source Verification & Reverse Searches

Reverse image and video search tools — including Google Images, TinEye, and InVID — can identify whether an image or video clip has appeared previously online in a different context. This is particularly effective when attackers repurpose existing public media rather than generating entirely new content.

Digital Signatures & Content Credentials

Digital signing of media at the point of capture — using camera firmware or authenticated recording applications — creates a cryptographic binding between the content and its origin device. Tampered content will fail signature verification, providing strong evidence of post-capture modification.

C2PA & Media Provenance Standards

The Coalition for Content Provenance and Authenticity (C2PA) defines an open technical standard for attaching verifiable provenance metadata — including capture device, creation software, and modification history — to media files. C2PA-compatible tools allow recipients to inspect a media file's full creation and editing chain.

Adoption of C2PA is growing among camera manufacturers, media organisations, and software providers. However, it is not yet universal, and significant quantities of legitimate media are distributed without provenance credentials.

> ⚠ The absence of C2PA credentials or digital signatures does not prove that media is synthetic. Much legitimate media lacks these credentials. Provenance analysis should inform, not determine, a deepfake assessment.

Key Limitations of Forensic Approaches

Metadata Can Be Stripped or Fabricated

Many social media platforms strip metadata on upload. Attackers can also forge metadata fields before distribution, making absent or anomalous metadata unreliable as a standalone indicator.

Provenance Standards Are Not Universal

C2PA adoption is voluntary and uneven. The majority of consumer devices and consumer-grade software do not yet embed C2PA credentials, meaning that most real media circulates without provenance attestation.

Re-encoding Destroys Forensic Signals

Transmitting media through messaging platforms, video conferencing tools, or social media re-encodes the content — overwriting compression signatures and destroying many forensic artefacts that would otherwise indicate synthetic origin.

Identity Verification and Authentication Controls

Technical detection tools are valuable, but procedural controls represent the most reliable safeguard against deepfake-enabled fraud. When an AI detection tool is unavailable, inconclusive, or simply not deployed in the communication channel being exploited, strong identity verification procedures provide the last line of defence. The core principle is straightforward: never authorise high-risk actions based solely on identity established through a single channel that could be spoofed.

1 Independent Out-of-Band Verification

When a request arrives via email, messaging app, or video call, verify the requester's identity through a separate, pre-established communication channel. Call the person back on a known phone number — not one provided in the request. Send a verification message to their official corporate email, not a reply to the original thread.

2 Multi-Factor Authentication for Sensitive Actions

Require MFA for all high-privilege actions — particularly password resets, account recovery, payment authorisation, and configuration changes. MFA should require a physical factor (hardware token or authenticator app) rather than relying solely on SMS, which is vulnerable to SIM-swap attacks.

3 Approval Workflows for Financial Transactions

Implement dual-authorisation or multi-party approval requirements for financial transfers above defined thresholds. Critically, approval workflows should require independent verification from each approver — not sequential approvals within the same communication thread that could all be compromised.

4 Acknowledge Limitations of Voice & Video Verification

Formally communicate to all relevant staff that voice recognition and facial recognition over phone or video call are not reliable identity verification methods. Establish this as an organisational policy, not merely a technical recommendation, to prevent social engineering that exploits employees' reluctance to challenge authority figures.

Sample Verification Workflow: Executive Payment Request

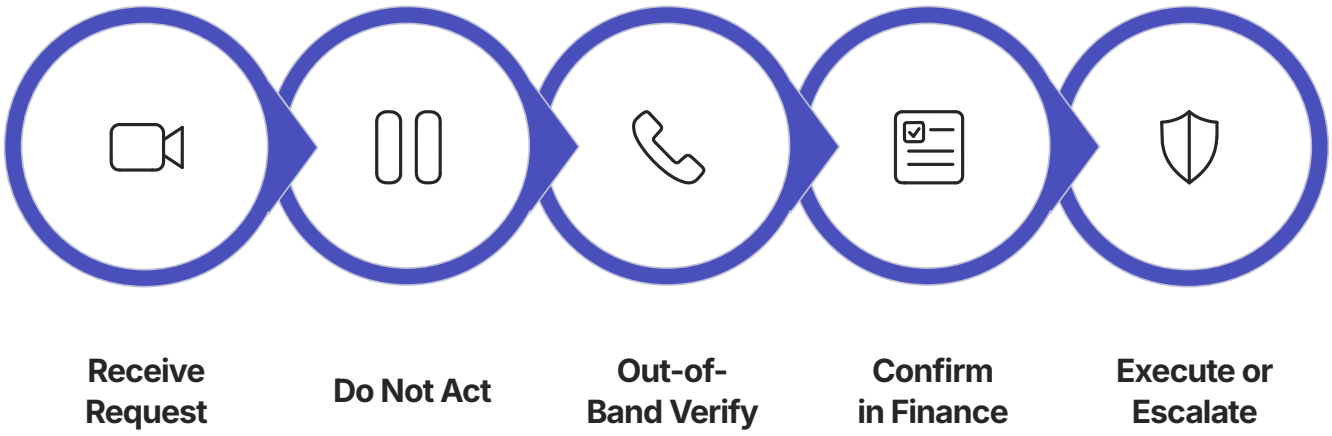

Preventing Deepfake Attacks in Organisations

Effective deepfake prevention requires a layered organisational approach that addresses people, processes, and technology simultaneously. No single control is sufficient — a comprehensive strategy combines identity verification procedures, process controls, employee awareness, and monitoring to reduce both the likelihood and impact of successful deepfake-enabled attacks.



Identity Verification Procedures

Define and document formal identity verification requirements for all sensitive actions. Specify which actions require out-of-band verification, which require dual authorisation, and what constitutes an acceptable verification method. Distribute these procedures to all relevant staff and enforce them consistently.



Protect Helpdesk & Account Recovery

Helpdesk and account recovery processes are high-value targets for deepfake-based social engineering. Require multiple verification factors for all account recovery requests. Never rely on voice recognition alone. Log all recovery requests and flag unusual patterns — such as multiple failed verifications — for review.



Financial Transaction Controls

Implement transaction approval thresholds with mandatory multi-party sign-off. Establish a standard that any change to payment details — bank account numbers, payment destinations — requires written authorisation through the official finance system, not via phone, video, or messaging app alone.



Employee Awareness Training

Train employees specifically on deepfake threats — not just phishing. Staff should understand that video calls and voice messages can be faked, that urgency and authority are common manipulation tactics, and that following verification procedures is not an act of distrust but a security requirement. Training should include practical scenario exercises.



Monitor for Unusual Activity

Deploy monitoring for unusual account activity, login anomalies, unexpected payment instructions, and out-of-pattern communication requests. Automated alerts for high-value transaction requests that deviate from established patterns can intercept attacks before execution.



Define Escalation Procedures

Establish clear escalation paths for suspected deepfake attempts. Employees should know exactly who to contact — their direct manager, the security team, or a fraud response function — and understand that reporting a suspicion, even if unfounded, is always the right action. Remove any cultural barriers to raising concerns about senior-figure requests.

Deepfake Incident Response Playbook

When a suspected deepfake attack is identified — whether a fraudulent transaction has already been initiated or the attempt was caught in time — a structured response process ensures that the incident is contained, evidence is preserved, and organisational defences are strengthened. The following playbook provides a practical step-by-step framework.

Step 1: Stop the Requested Action

Immediately halt any transaction, account change, or other action that is pending as a result of the suspicious request. If a payment has already been initiated, contact the financial institution immediately to request a recall or hold. Do not destroy or delete any communication records.

Step 2: Verify Identity Through Independent Channel

Contact the purported requester through a pre-established, independent communication channel — their known direct telephone number or official corporate email, not a contact method provided during the suspicious interaction. Confirm whether the request is legitimate.

Step 3: Preserve All Evidence

Preserve copies of the original audio, video, messages, email threads, call logs, and any associated files. Take screenshots with timestamps. Do not modify or forward the evidence. Store it in a designated secure location accessible to the incident response team.

Step 4: Investigate Associated Accounts & Channels

Review the communication channels used in the attack for signs of compromise — email account access logs, messaging platform sessions, video conferencing records. Identify any accounts that may have been accessed or used as part of the attack chain.

Step 5: Secure Affected Accounts

Reset credentials for any accounts identified as potentially compromised. Revoke active sessions and tokens. Apply additional authentication requirements where appropriate. Notify account owners through verified channels, not through the potentially compromised medium.

Step 6: Escalate to Security & Fraud Team

Notify the security team, fraud response function, and — where financial loss has occurred or is threatened — legal and compliance stakeholders. If the attack involves an apparent attempt to impersonate a named individual within the organisation, notify that individual as well.

Step 7: Document Findings & Update Controls

Produce a formal incident record capturing the attack timeline, evidence collected, verification steps taken, and outcome. Use findings to review and update prevention controls — identifying which process or technical control the attack exploited and what changes would prevent recurrence.

Practical Case Study: Deepfake-Based Executive Fraud

The following is a hypothetical scenario constructed to illustrate how a deepfake-enabled executive fraud attack might unfold, and how strong procedural controls can prevent financial loss.

How the Attack Begins

A finance team member receives an unexpected video call via a third-party messaging app from what appears to be their company's Chief Financial Officer. The caller's face, voice, and mannerisms are convincing. The CFO explains that a confidential acquisition is at a critical stage and that an urgent transfer of funds to a new external account must be completed within the hour to avoid losing the deal. The request emphasises confidentiality — the employee is asked not to discuss it with colleagues.

The attacker has constructed the synthetic video using publicly available footage of the CFO from an investor presentation. The voice was cloned using audio from the same recording. The messaging app number was registered using a temporary SIM.

Warning Signs Present

- Contact via a personal messaging app rather than official communication channels
- Unusual urgency and instruction to bypass normal approval processes
- Request for confidentiality — isolating the target from colleagues
- Payment to a new, unverified external account
- Subtle lip-sync latency and unnatural eye movement visible on close inspection

How Verification Prevents the Transfer

The employee, having received deepfake awareness training, recognises the combination of urgency, channel anomaly, and confidentiality instruction as social engineering indicators. Rather than initiating the transfer, they acknowledge the request and state they must follow standard financial approval procedures. They call the CFO's known direct office line. The CFO confirms no such request was made.

Investigation & Lessons Learned

Security Team Response

The employee reports the incident immediately. The security team preserves the video call recording, the messaging app conversation, and associated metadata. The synthetic video is submitted to a deepfake detection tool, which returns a high synthetic probability score. Investigators identify the source public videos used to train the model.

Documentation & Controls Update

The incident is formally documented. The security team issues updated guidance prohibiting financial authorisations via personal messaging platforms and reinforcing out-of-band verification requirements. A targeted deepfake awareness module is added to the finance team's mandatory training programme.

Key Lessons

The attack failed not because of AI detection software, but because an aware employee followed established verification procedures. This case reinforces that procedural controls — applied consistently — are the most reliable safeguard at the point of attack, while technical tools provide valuable supplementary evidence.

Deepfake Security Assessment Checklist

Use this checklist to assess your organisation's current readiness to detect, prevent, and respond to deepfake-enabled attacks. Each item represents a recommended control. Areas where controls are absent or partial should be prioritised for remediation.

Identity Verification Procedures

- [ ] Formal identity verification policy is documented and distributed
- [ ] Out-of-band verification is required for all high-value requests
- [ ] Voice and video alone are not accepted as sole identity verification
- [ ] Verification procedures cover third-party and vendor interactions

Financial Approval Controls

- [ ] Transaction approval thresholds are defined and enforced
- [ ] Dual authorisation is required above defined financial limits
- [ ] Payment destination changes require written sign-off via official systems
- [ ] Finance staff are trained on executive impersonation fraud patterns

Helpdesk & Account Recovery Security

- [ ] Account recovery requires multiple independent verification factors
- [ ] Voice recognition alone is not accepted for account recovery
- [ ] Recovery requests are logged and anomalies trigger review
- [ ] Helpdesk staff receive dedicated social engineering awareness training

Employee Awareness Training

- [ ] Deepfake-specific awareness training is delivered to all staff
- [ ] Training includes practical scenario exercises and indicators
- [ ] Staff know how to report suspected deepfake attempts
- [ ] Training is updated when new attack techniques emerge

AI-Based Media Detection

- [ ] Deepfake detection tools are evaluated or deployed for key channels
- [ ] Detection outputs are treated as probabilistic, not definitive
- [ ] Detection tool outputs are integrated into investigation workflows
- [ ] Limitations of detection tools are communicated to relevant teams

Evidence Preservation & Incident Reporting

- [ ] Incident response playbook covers deepfake-enabled attacks
- [ ] Evidence preservation procedures are defined and practised
- [ ] Secure storage for incident evidence is designated and accessible
- [ ] Reporting paths are clear and free of cultural barriers

Monitoring & Periodic Security Reviews

- [ ] Unusual account and transaction activity triggers automated alerts
- [ ] Communication channel anomalies are monitored and flagged
- [ ] Deepfake prevention controls are reviewed at least annually
- [ ] Post-incident reviews feed back into prevention control updates

This checklist is a practical self-assessment tool. Items marked incomplete represent recommended areas for control improvement. Prioritise gaps in identity verification and financial approval controls first.

Conclusion & Next Steps

Deepfake technology represents a genuine and evolving threat to organisational security. The attacks documented in this guide — executive impersonation, voice cloning, video conference fraud, and helpdesk manipulation — are not speculative risks. They require immediate, practical responses that combine multiple layers of defence.

Effective deepfake defence is not achievable through any single tool or technique. AI-based detection classifiers can identify synthetic media but produce probabilistic scores, not definitive verdicts. Media forensics and provenance tools add valuable context but have significant limitations in environments where metadata is stripped or standards are not universally adopted. Identity verification procedures provide the most reliable protection at the point of attack — but only when consistently followed and reinforced through employee awareness training. Incident response capability ensures that when an attack does succeed, its impact is contained and future defences are strengthened.

Organisations should treat deepfake prevention as an ongoing programme rather than a one-time implementation. The underlying AI technology continues to advance, and attack techniques will evolve accordingly. Regular review of procedures, training programmes, and detection tooling — informed by emerging threat intelligence — is essential to maintaining effective defences.

Develop Your Generative AI Cybersecurity Knowledge with GSDC

GSDC Generative AI in Cybersecurity Certification

The GSDC Generative AI in Cybersecurity Certification is designed for cybersecurity professionals, IT practitioners, and security leaders who want to build structured knowledge of generative AI applications and their security implications — including deepfake-enabled attacks, AI-driven phishing, and responsible AI governance.

The certification programme covers:

- Generative AI applications and their cybersecurity implications
- Identifying deepfake, phishing, and AI-driven impersonation threats
- Exploring AI security risks, vulnerabilities, and attack surfaces
- AI threat assessment frameworks and security practices
- Responsible AI use, governance, and organisational policy

Ready to Advance Your AI Security Knowledge?

Equip yourself and your team with the structured knowledge needed to understand, assess, and counter AI-generated threats in modern enterprise environments.

Who Should Attend

Cybersecurity professionals, IT security teams, fraud prevention specialists, risk and compliance officers, and security architects seeking formal recognition of AI security knowledge.

Explore the GSDC Generative AI in Cyber...

https://www.gsdcouncil.org/