What Is Social Engineering? Definition, Attacks, Scams | Proofpoint US

Published: 2021-12-28 · Archived: 2026-04-29 02:03:38 UTC

Table of Contents

Social engineering is a manipulation technique that targets human judgment rather than technical vulnerabilities. Attackers use trust, urgency, or deception to prompt actions against a person’s own interests. This could include stealing credentials or causing a company to send money via a fake wire transfer. The technology used is becoming increasingly sophisticated and can affect more people across a single attack. Using AI allows social engineers to create legitimate-looking emails at scale and even fake audio/video of executives to add legitimacy to a request.

Emails are no longer the only way for attackers to gain access to systems. Attacks now also come via SMS/text message, phone calls, social media, and collaboration platforms such as Slack and Teams. For those who deal with security issues and fraud prevention, social engineering has become one of the most significant problems companies face, including phishing, business email compromise (BEC), account takeover, and financial fraud.

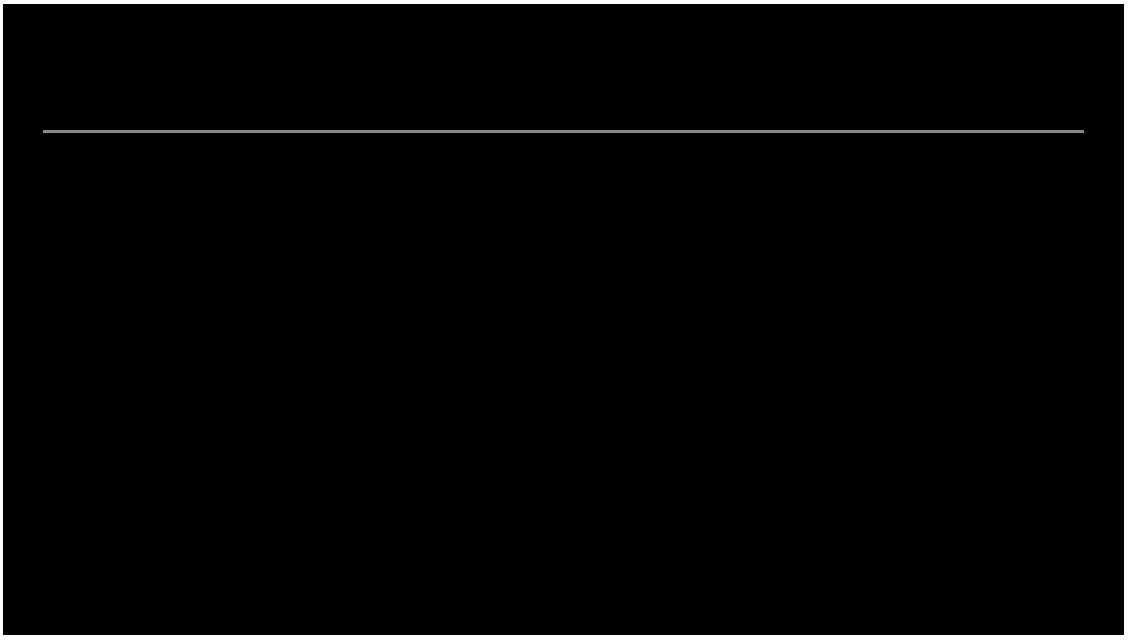

Here’s how your free trial works:

- Meet with our cybersecurity experts to assess your environment and identify your threat risk exposure
- Within 24 hours and minimal configuration, we'll deploy our solutions for 30 days
- Experience our technology in action!
- Receive report outlining your security vulnerabilities to help you take immediate action against cybersecurity attacks

Fill out this form to request a meeting with our cybersecurity experts.

Thank you for your submission.

Key Takeaways

- Social engineering uses psychological techniques, such as exploiting a person’s trust, sense of urgency, or authority, to prompt them to take unsafe or illegal actions. These exploits target people rather than systems.
- Any member of an organization could be targeted by social engineering attacks, from employees and executive staff to customers and vendors.
- Modern social engineering campaigns often use phishing and impersonation along with AI-generated content to make their messages seem more believable and difficult to identify.
- For many of today’s biggest security threats, social engineering provides the initial attack path. This includes BEC, account takeovers, identity theft, and various forms of fraud.
- To protect yourself from these types of threats, you need more than just technical tools. You need to include security awareness training, clearly defined verification processes, and multiple layers of defense.

Social Engineering Definition

Social engineering exploits psychological vulnerabilities to manipulate individuals into disclosing sensitive information or performing actions that compromise organizational security. While traditional attacks typically

target systems, social engineering targets people. This type of deception takes advantage of the way humans naturally react to trusting someone, recognizing authority, and being familiar with someone or something.

Today's attackers are now using AI in conjunction with deception and identity impersonation to make manipulation more difficult to detect. An example of this would be an attacker replicating an executive's voice to request that the company's finance department carry out fraud. The reason social engineering has proven to be such a successful method of attacking organizations is that it relies on behavior. The attackers take the time to study their intended victim(s), understand where they work and what they do, and develop their strategy accordingly.

How Does Social Engineering Work?

A threat actor might have a specific target in mind, or the attacker could cast a wide net to access as much private information as possible. Before a threat actor carries out a social engineering attack, their first step is to conduct due diligence on the targeted user or corporation. For example, the attacker could gather names and email addresses of the finance department staff from an organization's LinkedIn page to identify targeted victims and standard operating procedures.

The reconnaissance phase is critical to the success of a social engineering attack. The attacker must fully understand the business's organizational chart and target who has the authority to perform the actions necessary for success. In most attacks, social engineering involves the threat actor pretending to be someone the targeted user knows. The more information the threat actor collects about the targeted user, the more likely the social engineering attack will be successful.

With enough information gathered, the attacker can now carry out the next steps. Some social engineering attacks require patience to slowly build the targeted user's trust. Other attacks are quick, where the threat actor gains trust within a limited time by conveying a sense of urgency. For example, the attacker might call a targeted user and pretend to be an IT support staff member to trick the user into divulging their password.

What Are the Steps to a Successful Social Engineering Attack?

Just like most effective cyber-attacks, social engineering involves a specific strategy. Each step requires thoroughness because the attacker aims to trick the user into performing a particular action. Social engineering involves four steps. These steps are:

- **Information gathering**: This first step is critical to social engineering success. The attacker collects information from public sources like news clippings, LinkedIn, social media, and the targeted business website. This step familiarizes the attacker with the inner workings of the business departments and procedures.
- **Establish trust**: At this point, the attacker contacts the targeted user. This step requires conversation and convincing, so the attacker must be equipped to handle questions and persuade the targeted user to perform an action. The attacker must be friendly and might try to connect with the targeted user on a personal level.
- **Exploitation**: After the attacker tricks the targeted user into divulging information, exploitation begins. The exploit depends on the attacker's goals, but this step is when the attacker gets money, access to a system, steals files, or obtains trade secrets.

- **Execution**: With the sensitive information obtained, the attacker can now perform the final goal and exit the scam. The exit strategy includes methods to cover their tracks, including detection avoidance from the targeted organization's cybersecurity controls that could warn administrators that an employee had just been tricked.

Common Social Engineering Targets in Enterprises

Social engineering works best when attackers go after people with access, authority, or trust. In an enterprise environment, these targets make up more of the organization than most security teams realize.

- **Employees** represent the largest target group. Most phishing, smishing, and credential theft campaigns use broad targeting methods across the entire workforce.
- **Executives** are high-value targets for financial fraud, BEC, and deepfake impersonations. They are targeted by attackers to either obtain money from them or to impersonate them to get money from other employees.
- **Finance teams** are primary targets for wire fraud and invoice manipulation. If one employee approves a transaction, it could lead to significant financial losses for their employer.
- **Help desk/IT staff** are preyed upon through pretexting and impersonation to reset credentials, obtain new account access, or gain new system privileges.
- **Customer service teams** are vulnerable to account-takeover attacks, in which attackers impersonate legitimate customers to gain access to accounts or obtain customer personal data.
- **Third-party vendors and partners** can serve as an indirect point of entry into your systems because of the trusted relationships they've developed with your company.

How Social Engineering Has Evolved

Social engineering has been an issue for many years. These cyber-attacks began with simple phone scams and generic phishing emails that would cast a broad net in hopes of finding enough people to scam. The tactics were usually very basic and fairly easy to identify, but they had the potential to succeed because they were so widespread.

Things have changed significantly today. Modern social engineering campaigns are designed as multi-channel efforts. They can move across email, SMS, voice calls, social media, and even your work-based communication platforms. AI has accelerated the shift from manual, opportunistic scams to personalized, convincingly human automated attacks.

One of the greatest changes we have seen in modern social engineering is the level of precision. These attackers will do their homework to learn about you, mimic how someone familiar to you sounds and act/look like them, and create messages that seem legitimate. This isn't just a numbers game anymore. Modern social engineering is much more of a targeted campaign. As such, it is much harder to detect than it was in the past.

Social Engineering in the Age of AI

AI has transformed the potential for large-scale social engineering. Threat actors are now able to leverage large language models to produce custom, contextually relevant communications within seconds, at a price point that allows nearly all types of threat actors access to highly-targeted spear phishing. A study conducted in January of

2025 demonstrated that AI-powered spear phishing had a 54% click-through rate (CTR), which is comparable to that of a human expert, but at less than one-tenth the cost.

This attack vector extends far beyond email. Voice-cloning applications may mimic a CEO's tone and style based on a few minutes of audio. AI chatbots are able to maintain a believable conversation in near-real time, pretending to be IT support or a financial representative of a firm until their target complies. In the 2024 CFO deepfake video conference case, there were no real participants on the call, and the company lost $25 million before identifying the scam.

Common Signs of a Social Engineering Attack

Social engineering works because it triggers instinctive responses before the target has time to think critically. Urgency, fear, authority, and curiosity are the primary lenses used to induce those responses. Being aware of these signals is the first line of defense against them.

A few patterns show up consistently across modern attacks:

- **Pressure to act fast**: When you need to act fast with financial transfers, reset credentials, or approve access to a resource, that is usually a tell-tale sign of social engineering. Executives often receive high-pressure financial requests from what they believe to be their colleagues or vendors.
- **Identity impersonation**: Attackers pretend to be executives, IT staff, vendors, or customer Support agents. This happens via email, phone calls, Slack, Teams, and even text messages. Cloned voices and AI-generated audio are now making it harder to detect phone-based impersonations.
- **Polished, convincing language**: AI-generated messages no longer have typos and awkward phrasing, which made such exploits easy to spot in the past. If a message feels slightly off despite looking professional, that instinct is worth listening to.
- **Suspicious links and QR codes**: Phishing links embedded in messages are a long-standing tactic. QR code scams, sometimes called quishing, are a growing variation that bypasses many traditional security filters.
- **Account recovery and help desk requests**: Support teams are frequent targets. Attackers pose as employees who need urgent credential resets or account changes. Always verify identity through a separate channel before making any access changes.
- **Unverifiable senders**: If somebody is unwilling or incapable of confirming who they are, that alone should cause you to pause. This goes for every channel.

Examples of Social Engineering Techniques

The aim of social engineering is to deceive people into taking actions that they wouldn't ordinarily take. The techniques listed below are some of the most common and damaging techniques being employed today.

Digital Impersonation

Attackers use digital channels to impersonate someone recognizable. They can ask for your password, money, etc. Most of the time, these attacks come via email or other digital channels.

- **Phishing**: These deceptive emails are designed to obtain user credentials, distribute malware, or direct users to fraudulent sites. Phishing messages are becoming more elaborate thanks to AI, so it is becoming increasingly difficult to determine if the message is legitimate or not.
- **Smishing**: This is a form of phishing that uses SMS/text messages to send users malicious links or prompt users to contact an illegitimate phone number. Users typically respond to an SMS more quickly and with less scrutiny than to an email.
- **Vishing**: The attacker uses a voice to impersonate a legitimate caller, including an executive, a vendor, or an IT staff member. Due to advancements in voice-cloning technology, vishing has become much more believable.
- **QR phishing (quishing)**: Scammers embed malicious QR codes—in e-mails, PDFs, and physical signs—that lead users to websites where they are asked for their credentials. QR codes can bypass many of the typical link-scanning filters.
- **Deepfake fraud**: Deepfakes are generated using AI and include both audio and video. They can be used to impersonate a company executive, a colleague, or a vendor. These types of attacks are becoming more prevalent for use to make unauthorized transactions or influence employees during live conversations.
- **Collaboration app scams**: Scammers are now utilizing collaboration apps to impersonate executives or IT personnel. The scammers will create urgency around their request and may compromise an internal account to seem like a legitimate request.
- **Social media impersonation**: Threat actors are creating fake social media profiles that mimic executives, colleagues, or brands. Once the trust is built, the scammer will ask the target for sensitive information or ask the target to visit malicious content.

Access Manipulation

These tactics involve obtaining unauthorized access to systems through deception, persistence, or physically being present.

- **Pretexting**: Attackers create a fabricated story in order to gain the trust of the target. A classic example of pretexting is when a scammer poses as a bank representative to verify a customer's account information after a reported security breach.
- **Help desk deception**: Support teams are high-value targets. Attackers impersonate employees in need of urgent credential resets or account changes. Verification workflows that rely on a separate, established channel are the most effective defense.
- **MFA fatigue**: Threat actors continue to repeatedly prompt the target for multiple forms of authentication until the target becomes frustrated or confused enough to accept one. This tactic was used in some high-profile breaches and can be detected only through behavioral analysis.
- **Tailgating**: This exploit occurs when an unauthorized individual follows an authorized employee through a secured door. This allows the unauthorized individual to bypass physical security measures completely. Tailgating does not require technical expertise; it relies on the employee's courtesy and distraction.

Incentive and Lure Tactics

These techniques use promises, rewards, or fabricated opportunities to lower a target's guard.

- **Baiting**: Scammers lure victims into accessing a malicious link or downloading a piece of malware, which is often in the form of free software, a gift card, or exclusive content.
- **Quid pro quo**: These attacks involve sharing something of value in exchange for access to the target's system or sensitive information. Sometimes, quid pro quo is used by disgruntled employees who are contacted by outside threat actors.
- **Fake offers and surveys**: Disguised customer service interactions, prize notifications, or survey invitations are used by scammers to gather the target's personal information or to redirect the target to a website that collects credentials.

How to Prevent Social Engineering Attacks

Urgency is the most reliable signal that something deserves a second look. Attackers use it deliberately to compress the time between a request and a response. Slowing down is not a sign of inefficiency. It is the right instinct. A few practices that hold up across every channel:

- **Verify through a separate channel**: If a request arrives by email, confirm it by phone or in person before acting. If it arrives by phone, follow up through a known contact method. This applies to financial approvals, credential resets, and access changes. Executives receiving urgent wire transfer or approval requests should treat out-of-band verification as a standard step, not an exception.
- **Do not trust urgency alone**: Pressure to act immediately, even from a familiar name or voice, is reason to pause. Legitimate requests from colleagues, vendors, and executives can almost always wait a few minutes for verification.
- **Treat AI-generated content as a credibility challenge**: Polished language, familiar tone, and accurate context no longer indicate a message is real. AI makes impersonation more convincing across email, SMS, and collaboration platforms. Apply the same scrutiny to a well-written message as a poorly written one.
- **Apply the same skepticism to collaboration tools as email**: Impersonation in Slack, Teams, and similar platforms is a growing tactic. A message from a colleague's account requesting urgent action should be verified the same way an email would be.
- **Be cautious with voice and video requests**: Voice-cloning and deepfake technology have made audio and video less reliable as identity proof. If a call or video request involves a sensitive action, verify through a second channel before proceeding.
- **Navigate directly to sites rather than clicking links**: If a message claims to be from a known service or vendor, open a browser and go directly to the site. The same applies to QR codes in unsolicited messages or physical materials.

Help desk and SOC teams require strong identity verification before any credential reset or access change. An attacker's best path through your organization may be a single convincing phone call.

Why Social Engineering Remains So Effective

Social engineering defense techniques have advanced dramatically. However, these threats remain the leading threat to organizations because the most difficult defense to patch is human judgment. These include urgency, authority, and familiarity, among others. As such, attackers are continually improving their ability to manufacture these elements for use against individuals.

What has changed is scale and precision. Attackers can now use AI to generate a high volume of personalized, contextually relevant messages, which was impossible with previous technology. 2025 research found that AI-generated phishing attacks outperformed elite human red teams by 24%, a meaningful shift from just two years earlier, when AI lagged significantly behind.

The attack surface has also widened. Threats no longer arrive through a single channel or identity. They move across email, voice, SMS, and collaboration platforms in coordinated sequences. According to the Verizon 2025 DBIR, the human element was a factor in roughly 60% of breaches. That number has held steady for years, and that consistency is the point.

Emerging Trends in Social Engineering

Deepfakes have gone from a hypothetical risk to an active threat, with hackers leveraging AI-generated audio and video to impersonate executive-level personnel, vendors, and colleagues. The quality of deepfakes is such that voice authentication can no longer be relied upon to identify a person.

The number of malicious QR code detections reported by Kaspersky increased fivefold between August and November of 2025. This type of attack is so successful at getting past people's defenses because it often bypasses email filters and sends users to credential-harvesting sites on their mobile devices, which typically have less robust security software than desktops.

Collaborative workspaces have emerged as a major attack vector. There are increasing instances of impersonation on these platforms, especially when an attacker compromises a legitimate internal account, giving the impersonator some credibility.

There is growing concern among security professionals about AI-enabled autonomous capabilities for executing social engineering campaigns. Security teams are already planning how to deal with autonomous AI agents that can design and execute all aspects of a social engineering campaign, from research to lure development to delivery via any medium (e.g., email, phone).

FAQs for Social Engineering

What is social engineering in simple words?

How is social engineering different from phishing?

How are attackers using AI in social engineering?

Can deepfakes be used in social engineering attacks?

How can organizations prevent social engineering?

Why is social engineering effective even with strong security tools?

What is the most common form of social engineering?

Is social engineering illegal?

How common is social engineering?

What is social engineering in simple words?

Most people think of cyber-threats as malware or a hacker exploiting vulnerabilities in software. However, social engineering is a threat where an attacker tricks a targeted user into divulging sensitive information by pretending to be a familiar person or service. The attacker might trick a targeted user into divulging their password, or the attacker could trick the targeted user into sending money by pretending to be a high-level executive. Attackers' goals in a social engineering campaign vary, but generally, the attacker wants access to accounts or to steal the user's private information.

How is social engineering different from phishing?

Phishing is one type of social engineering, but social engineering is the broader category. It includes any technique that manipulates people through deception, whether over email, phone, SMS, video, or in person. Phishing specifically refers to deceptive messages designed to steal credentials or deliver malware. Social engineering also encompasses vishing, pretexting, deepfake fraud, help desk deception, and more.

How are attackers using AI in social engineering?

AI has increased the speed and scalability of social engineering. AI uses large language models to produce high-quality, personalized, and professional-looking emails in bulk. This eliminates some of the characteristics that make phishing so easy to identify. The use of AI also enables the development of voice cloning, synthetic videos, and automated reconnaissance that can provide an attacker with information about their intended victim.

Can deepfakes be used in social engineering attacks?

Yes, they are being used. An example of how this has been done was seen in a very publicized incident involving a British engineering firm called Arup. They were targeted with an AI-generated deepfake video. The attackers pretended to be the company's CFO during a video conference and caused the firm to lose roughly $25 million. Between the first and second half of 2024, voice-cloned attacks went up 442%. Deepfakes are no longer a limited tool for well-funded and sophisticated attackers. They are now being used in both financial fraud and executive impersonation schemes.

How can organizations prevent social engineering?

Organizations need to take a multi-layered approach. This includes continuing to train employees on security awareness, having out-of-band workflows for sensitive requests, implementing strong identity controls, and using behavior-based detection to monitor and alert when an employee performs an abnormal action. Organizations also need to document how the identity of callers to the helpdesk and SOC teams will be verified prior to resetting any credentials or making any changes to access. Culture is important as well. If employees believe they can safely report suspicious activities, then they become a significant deterrent against social engineering.

Why is social engineering effective even with strong security tools?

Impersonations bypass security technologies because they target individuals directly. A sophisticated impersonation doesn't have to go through a firewall to reach an individual. Or evade an endpoint agent. It simply has to get someone to act on the request. Most security tools can identify and block known types of malicious content. However, they are less likely to identify a convincing request from someone who appears to be a co-worker or supervisor. The most obvious characteristics of a malicious message that would alert a user to its potential harm are removed by artificial intelligence. This places nearly all of the burden for determining whether a communication is legitimate on the individual (which is exactly where an attacker wants it).

What is the most common form of social engineering?

The term "social engineering" is a broad term that covers many cyber-criminal strategies. Social engineering involves human error, so attackers target insiders. The most common form of social engineering is phishing, which uses email messages. Under the umbrella of phishing are vishing (voice) and smishing (text messages). In a typical phishing attack, the goal is to obtain information for monetary gain or data theft.

In a phishing email, the attacker pretends to be a person from a legitimate organization or a family member. The message might ask for a simple reply, or it will contain a link to a malicious website. Phishing campaigns can target specific people within an organization – spear phishing – or the attacker can send hundreds of emails to random users, hoping that at least one falls for the fraudulent message. Untargeted phishing campaigns have a low success rate, but it doesn't take many successful messages for an attacker to obtain the necessary information for monetary gain.

The two phishing variants – smishing and vishing – have the same goals as a general phishing campaign but different methods. A "smishing" attack uses text messages to tell targeted users that they have won a prize and need to pay a shipping fee to receive their gifts. "Voice" phishing requires voice-changing software to trick users into thinking the attacker is someone from a legitimate organization.

Is social engineering illegal?

Yes, social engineering is illegal because it uses deception to trick individuals into revealing sensitive information or granting access to systems. These attacks often lead to more serious crimes, including fraud, identity theft, and unauthorized access to networks or financial accounts.

A common example is consumer fraud, where attackers impersonate trusted organizations to request financial details or payments. Once obtained, this information can be used to steal money or sold on illicit markets.

Penalties vary depending on the scale and impact of the attack. Smaller offenses may result in fines or short-term jail time, while larger or repeated attacks can lead to significant prison sentences, higher fines, and civil lawsuits brought by affected victims.

How common is social engineering?

Very common, and growing. According to the 2025 SANS Security Awareness Report, 80% of organizations rank social engineering as their top human-related risk. Palo Alto Networks' Unit 42 found that social engineering was

the leading initial access vector in incident response cases between mid-2024 and mid-2025, accounting for 36% of all incidents. In two-thirds of those cases, attackers targeted privileged or executive accounts.

The goals behind these attacks have also broadened. Credential theft remains the most common outcome, but social engineering now drives financial fraud, account takeover, data theft, and ransomware deployment. Business email compromise alone generated $2.8 billion in reported losses in 2024, according to the FBI.

Take Ownership of Your Data with Proofpoint

Companies that invest in data security and governance are better able to control where sensitive information is stored, who can access it, and how it moves throughout their environment. To protect data effectively, you need to do more than just stop threats at the perimeter. It requires ongoing visibility into insider behavior, unauthorized access patterns, data governance policies, and internal systems that can adapt as data moves. When securing and preventing data loss is a top priority, the right mix of discovery, classification, and access controls can help businesses stay ahead of both intentional misuse and unintentional exposure.

See why enterprises trust Proofpoint for comprehensive data protection that addresses tomorrow's threats. Contact Proofpoint today.

Related Resources

The latest news and updates from Proofpoint, delivered to your inbox.

Sign up to receive news and other stories from Proofpoint. Your information will be used in accordance with Proofpoint's privacy policy. You may opt out at any time.
