GROUP-IB

White Paper

The Voice of Fraud

Deepfake Vishing and the New Age of Social Engineering

GROUP-IB.COM

Table of contents

Disclaimer 03

Acknowledgements 03

Introduction 04

Trust is the New Attack Surface 05

- Abuse of legitimate AI voice cloning platforms in the wild 06

AI voice cloning: anyone’s voice, at anyone’s fingertips 06

- Lower barriers to entry 06
- High fidelity cloned voices 08
- Voice cloning and differentiation experiment with CNA 08

Caller-ID spoofing: How telecoms infrastructure can be compromised 09

- Understanding caller ID 09
- How caller ID data is transmitted across different telephony systems 10
- Three levels of caller ID spoofing in the telephony ecosystem 11
- How fraudsters conduct VoIP/full SIP spoofing 12

Organizational blind spots: voice = identity in corporate environments 17

Detection and defense: strategies across stakeholders 18

- Telecom sector: strengthening the infrastructure layer 18
- Corporate environments: zero trust for voice-only approvals 19
- Individual users: awareness and personal safeguards 19

Conclusion 20

Disclaimer

01

The report is for information purposes only and Group-IB is limiting its distribution. Readers are not authorized to use it for commercial purposes or any other purposes not related to training or personal non-commercial use. Group-IB grants readers the right to use the report worldwide by downloading, reviewing, and quoting it to the extent justified by legitimate citation, provided that the report itself (including a link to the copyright holder's website on which it is published) is given as the source of the quote.

02

The entire report is subject to copyright and protected by applicable intellectual property law. It is prohibited to copy, distribute (including by placing on websites), or use any of its content without the copyright holder's prior written consent.

03

In case of copyright infringement, Group-IB will have the right to approach a court or other state institution to protect its rights and interests and seek punishment for the offender as provided by law, including recovery of damages.

Acknowledgements

Authors

Yuan Huang
Senior Fraud Analyst,
Asia-Pacific

Group-IB would also like to thank CNA (Channel NewsAsia) for their participation in the AI deepfake vishing experiment, as well as Group-IB's Fraud Protection team for its contributions to this report and experiment.

Introduction

In recent years, AI-enabled fraud has seen an upward trend globally, specifically those involving the use of deepfaked voice, images and videos. According to industry projections, global financial losses attributed to AI-enabled fraud are expected to reach US$40 billion by 2027, up from approximately US$12 billion in 2023.

This report seeks to highlight the growing and increasingly sophisticated threat posed by AI-powered voice impersonation attacks, also known as deepfake vishing. These attacks exploit critical and long-standing vulnerabilities within telecommunications infrastructures and organizational trust models. Specifically, this report examines how caller-ID spoofing and a continued reliance on voice-based identity verification have created blind spots that are being actively exploited by threat actors.

Two key trust signals are being manipulated: the phone number shown on the recipient's screen, and the familiar voice of someone with authority. By cloning the voice of a senior executive-level figure — such as a CFO or department head — and spoofing their caller ID, attackers use highly convincing scenarios that prompt urgent financial transfers or disclosure of sensitive information.

This threat is particularly urgent due to its **scalability and stealth:**

- **AI voice cloning tools** are inexpensive, publicly accessible, and require minimal technical expertise.
- **Telecommunications networks** lack robust caller-ID authentication, making spoofing straightforward to execute.
- **Organizations continue to depend on voice-only verification**, exposing them to impersonation risks.
- **Psychological manipulation techniques**, such as urgency, fear, and appeals to authority, are effectively leveraged to coerce action.

This report demonstrates how systemic vulnerabilities intersect with rapid advancements in generative AI and behavioral exploitation tactics. In doing so, it emphasizes the urgent need for telecom providers, corporations, and professionals to modernize verification protocols — before such attacks become a normalized and highly damaging component of the threat landscape.

Trust is the New Attack Surface

In an era where any voice can be cloned and any number can be spoofed, trust in familiar signals is being weaponized. AI-generated voices, paired with caller-ID spoofing, are enabling highly targeted fraud. The implications extend far beyond individual deception: executives, financial teams, and public officials are now routinely being targeted through this advanced form of impersonation.

This report explores the rise of AI-driven vishing attacks in conjunction with caller-ID spoofing, focusing on the convergence of:

- **Voice cloning technology,** which enables precise impersonation of individuals.
- **Telecom protocol weaknesses,** which permit spoofing of trusted numbers.
- **Organizational dependence on voice-based verification,** which increases exposure to deception

By dissecting the structural enablers of this threat, this report aims to help organizations foresee and fortify against a new category of fraud — one that undermines not only technical systems but interpersonal trust itself. In this landscape, verification must be redefined. Voices and caller IDs can no longer be assumed genuine without supporting validation.

Abuse of legitimate AI voice cloning platforms in the wild

The operational deployment of AI voice impersonation in fraud schemes has already been observed in-the-wild.

United Kingdom

In one high-profile incident, a UK-based energy firm suffered a financial loss of **$243,000** after receiving a call from what appeared to be their CEO. The voice, generated using AI, convincingly mimicked the German executive and was used to urgently request a wire transfer.

Hong Kong

In February 2025, it was reported that a merchant in Hong Kong lost **HK$145-million** (approximately US$18.5-million) in USDT (stablecoin) after scammers used AI to clone the voice of the financial manager of the mainland firm.

In these two examples, although the caller's number was unknown, no immediate suspicion was aroused due to the authenticity of the voice. Had the phone number been spoofed as well, the likelihood of detection would have diminished even further, particularly in high-pressure contexts.

AI voice cloning: anyone's voice, at anyone's fingertips

Lower barriers to entry

Voice cloning technology has rapidly evolved to the point where generating convincing imitations of real individuals requires minimal effort, expertise, or cost. With as little as a few seconds of audio — often harvested from publicly available recordings on social media, or purchased through underground marketplaces — threat actors are now able to synthesize highly realistic voice clones of CEOs, financial executives, or government officials.

Commercial voice synthesis platforms such as **ElevenLabs, Resemble.ai,** and **PlayHT** provide advanced cloning capabilities through user-friendly web portals and API integrations. In parallel, open-source toolkits like Coqui TTS and Real-Time Voice Cloning are readily accessible on GitHub under permissive licenses. These tools typically require no more than a text input to produce speech in a cloned voice. Many platforms offer free trials or affordable subscription tiers, allowing high-quality voice impersonation to be carried out at negligible financial cost — often for less than the price of a cup of coffee.

At this point we have to emphasize that these platforms were originally created for legitimate purposes. However, fraudsters may misuse or abuse these platforms, in order to carry out vishing attacks.

Companies	What customer information is required to create a custom voice clone	Entry-level pricing for custom voice cloning	Whether there are technological barriers to generating a clone (e.g., training data amount, approval steps)	Whether their privacy policy prohibits malicious use
ElevenLabs	First name, email, credit-card (plus ≥30 min voice sample)	Starter plan ~$5/mo (includes “Instant Voice Clone”)	None (just a checkbox confirmation)	Terms prohibit illegal/malicious use
Resemble.ai	Name, email, payment info (must own or have rights to uploaded voice)	Creator plan $19/mo (15,000 sec, 1 pro clone) (or pay-as-you-go ~$0.018/sec)	None (just a checkbox confirmation)	TOS forbid malicious, deceptive or harmful uses
PlayHT	Name and email (standard signup)	Free tier with 1 instant clone (12.5k chars) Professional plan ~$29.25/mo (voice cloning enabled)	None (just a checkbox confirmation)	Terms forbid harassment, impersonation, and fraud
Speechify	Email (account signup)	Free (voice cloning available without extra charge)	None (just self-attested consent)	Terms forbid illegal or fraudulent uses
Descript	First name, email, credit-card (plus ≥30 min voice sample)	Starter plan ~$5/mo (includes “Instant Voice Clone”)	None beyond consent (checkbox after recording)	Requires explicit speaker consent for cloning
LOVO	Name and email (account signup)	Basic plan $24/mo (2 hr gen, up to 5 custom clones)	None (just a checkbox confirmation)	TOS forbid illegal, harmful or defamatory use

High fidelity cloned voices

Modern AI voice models are designed to reproduce far more than basic vocal characteristics. In addition to pitch and accent, these systems are trained to emulate natural speech patterns, including rhythm, pauses, emotional tone, and subtle non-verbal cues such as sighs, breathing, and vocal fry. As a result, the synthetic output closely mirrors the unique vocal identity of the target individual.

Empirical research underscores the effectiveness of these models. A 2024 study conducted by the University of California found that participants were unable to reliably distinguish between real and AI-generated voices. The cloned voice was mistaken for the real speaker in nearly 80% of test cases, while AI-generated samples were correctly identified only 60% of the time—comparable to random guessing. These findings highlight the growing indistinguishability of synthetic voices, even by attentive human listeners.

Voice cloning and differentiation experiment with CNA

Group-IB, together with Singapore-based broadcaster CNA (Channel NewsAsia), carried out an experiment to show how legitimate AI-powered voice cloning platforms — widely available online, inexpensive, and easy to use — could be exploited by cybercriminals for vishing attacks.

Figure 1.
Yuan Huang, Group-IB's Senior Fraud Analyst records the voice sample of CNA journalist Natasha Ganesan during the experiment. Image credit: CNA

In an article published about her experience during the vishing experiment, experiment, the journalist noted that:

- A voice recording of just 10 seconds would be sufficient for the voice-cloning platform to generate a cloned voice.
- Subscription fees for such platforms range from US$3 to US$10 a month.
- The user-friendly interface does not require any technical knowledge in voice cloning.
- While not entirely flawless, the cloned voice shared similarities to the journalist's actual speech patterns, including the habitual pause between words.

Although all the recipients — who had been informed in advance to expect both a real and a cloned voice call from the journalist — were able to distinguish between the AI-generated voice and the genuine one, experts interviewed by CNA noted that voice communication can still be highly persuasive in social engineering, particularly when it is used to create a sense of urgency or trigger an emotional response.

The entire experiment was also filmed by CNA, highlighting the affordability, ease-of-use, and speed of publicly available AI voice cloning platforms today. Watch how the experiment was conducted via the official CNA website.

Caller-ID spoofing: How telecoms infrastructure can be compromised

Understanding caller ID

Caller ID is a standard telecommunications feature intended to display the caller's phone number — and in some cases, their name or business — on the recipient's device. While its primary purpose is to help users identify incoming calls, this system is inherently vulnerable to **caller ID spoofing**: a technique in which the caller falsifies the displayed number to appear as a trusted individual or entity.

Caller ID information is set by the caller's phone system at the time of the call. Although telecom providers and devices may perform additional lookups (e.g., to match a number with a known business or geographic location), the core phone number shown on the recipient's screen is **provided by the caller**. This fundamental design flaw allows malicious actors to easily manipulate caller ID values with minimal technical effort.

How caller ID data is transmitted across different telephony systems

Plain old telephone system (POTS)

In traditional analog landline networks, caller ID data is delivered using **Frequency Shift Keying (FSK)** — a modulation method that encodes digital data into audio frequencies. This FSK signal is transmitted between the first and second ring of an incoming call and contains the caller's number, along with optional metadata such as time and date. The receiving phone then decodes this signal to display the caller ID.

Personal home phones and mobile phones

For both personal landlines and mobile networks, caller ID information is automatically injected by the telecom provider. End users have no direct control over this process. The telecom provider pulls or passes the caller ID information during call setup, and Caller ID reliability in this context depends largely on the upstream caller's system and the integrity of the telecom carrier.

Voice over IP (VoIP) systems

VoIP systems transmit calls over internet-based networks rather than legacy telephone infrastructure. These systems rely on the Session Initiation Protocol (SIP) for call signaling and management. When a VoIP call is placed, the caller ID is typically set in the "From" header of the SIP INVITE request — generated by the calling device (e.g., a softphone, SIP client, or VoIP phone).

```
INVITE sip:bob@biloxi.example.com SIP/2.0
Via: SIP/2.0/TCP client.atlanta.example.com:5060;branch=z9hG4bK74bf9
Max-Forwards: 70
From: Alice <sip:alice@atlanta.example.com>;tag=9fxced76sl
To: Bob <sip:bob@biloxi.example.com>
Call-ID: 384827629822018851l@atlanta.example.com
CSeq: 2 INVITE
Contact: <sip:alice@client.atlanta.example.com;transport=tcp>
Diversion: Carol <sip:carol@atlanta.example.com>;privacy=off;reason=no-
answer;counter=1;screen=no
Remote-Party-ID: Alice <sip:alice@atlanta.example.com>
P-Asserted-Identity: Alice <sip:alice@atlanta.example.com>
P-Charge-Info: <sip:eve@atlanta.example.com>
P-Source-Device: 216.3.128.12
Content-Type: application/sdp
Content-Length: 151
X-BroadWorks-DNC: network-address=sip:+9876543210@127.0.0.101;user=phone
User-Agent: X-Lite release 1104o stamp 56125 v=0 o=alice 2890844526 2890844526 IN
IP4 client.atlanta.example.com s=- c=IN IP4 192.0.2.101 t=0 0 m=audio 49172 RTP/AVP
0 a=rtpmap:0 PCMU/8000
```

Figure 2. An example of a SIP INVITE request.

The VoIP call flow

The SIP INVITE is first sent to a SIP server or a Private Branch Exchange (PBX), which may enforce caller ID policies (e.g., replacing user-supplied phone numbers with predefined ones for consistency or compliance). Businesses often use PBX systems to ensure all outbound calls display a consistent caller ID (e.g., main office number).

The INVITE is then forwarded to a **VoIP provider**, who acts as the gateway to the public telephone network. Depending on the provider's policy, the original caller ID may be preserved or overwritten.

Once validated, the SIP INVITE is routed through a **PSTN (Public Switched Telephone Network) gateway**, where SIP signaling is translated into legacy telecom protocols (e.g., FSK) for delivery to the recipient.

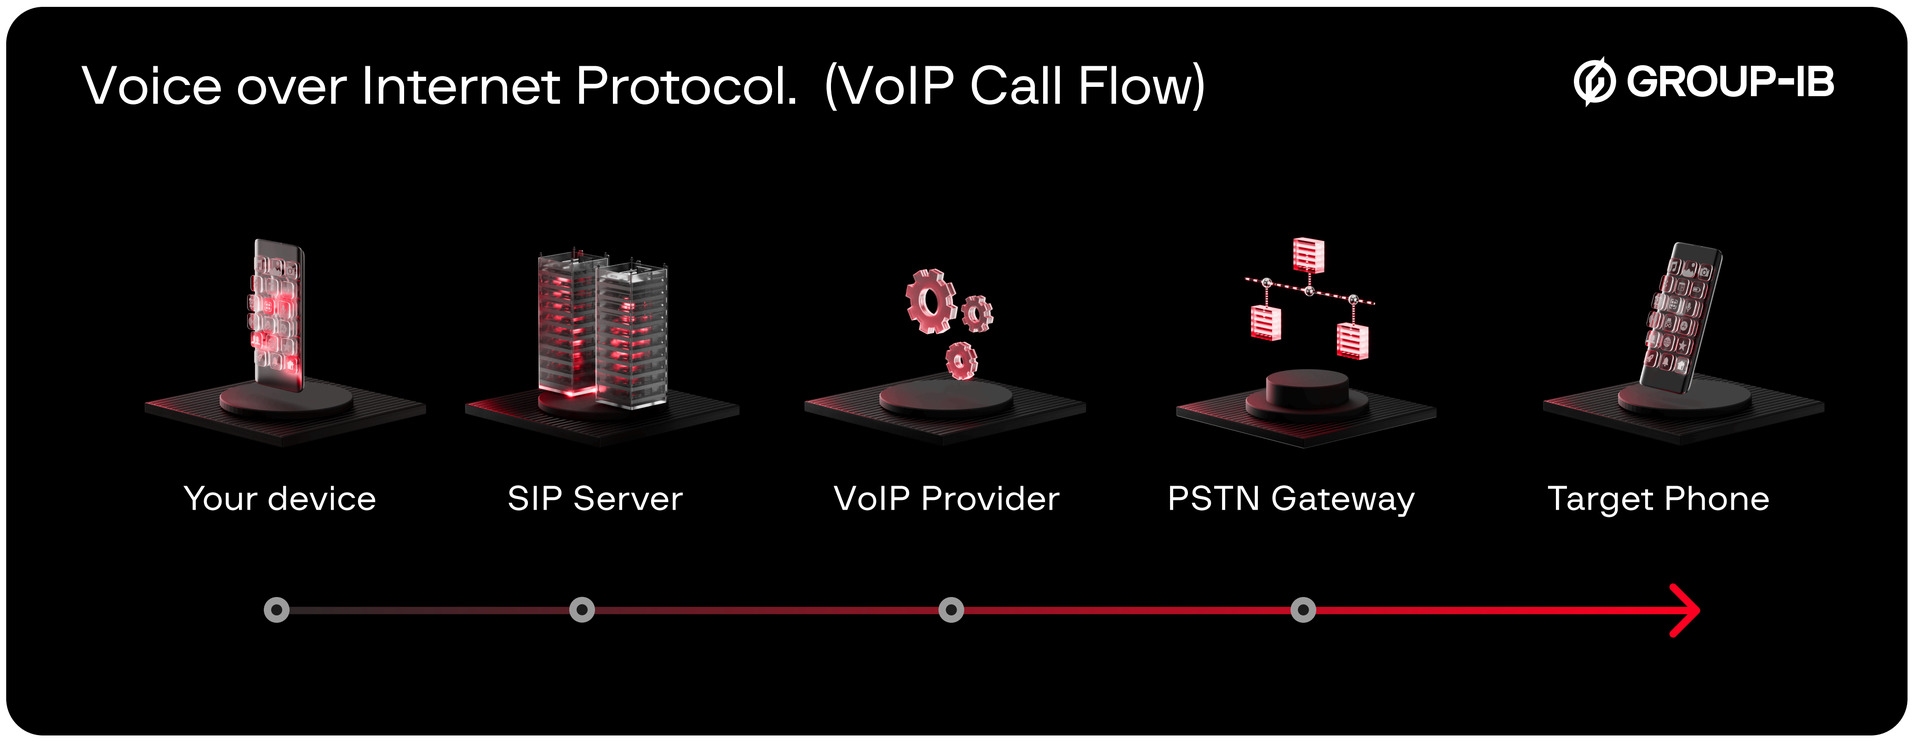

Figure 3.
An illustration of the Voice-over-IP (VOIP) workflow.

Three levels of caller ID spoofing in the telephony ecosystem

Carrier-Level Spoofing

This high-level spoofing targets traditional telecom networks, including POTS and cellular systems. It often involves manipulation of **SS7 (Signaling System No. 7)** protocols or unauthorized access to carrier infrastructure. Due to its technical complexity and high cost, this method is typically associated with insider threats or nation-state actors. When executed successfully, it is extremely difficult to detect and highly reliable in targeted attacks.

VoIP/Full SIP Spoofing

The most prevalent spoofing method today. Attackers either operate their own SIP infrastructure or rent SIP trunks, then craft SIP INVITE messages with forged "From" headers. The ability of this spoofing to succeed depends heavily on downstream VoIP providers and their enforcement of caller ID validation. If accepted and routed to PSTN gateways, the spoofed caller ID will be shown on the recipient's device. This method is frequently used in social engineering, robocalling, and financial fraud.

Display-Level Spoofing

This superficial form of spoofing is seen in applications like TextNow or Dingtone, where users manipulate the caller ID **at the UI level.** No legitimate SIP or telecom signaling is involved; spoofed numbers are only visible within the app ecosystem. These calls are often blocked by spam filters and are ineffective against enterprise systems or high-trust targets. This technique is typically used for pranks or casual anonymity.

How fraudsters conduct VoIP/full SIP spoofing

Modern businesses increasingly rely on VoIP (Voice over IP) phone systems due to their low cost and flexibility. This however, also makes VoIP systems a potential target for committing fraud. Three common **VoIP/Full SIP Spoofing** methods have been identified:

- **Manual Modification of SIP Headers via SIP Trunking Services**
- **Deployment of Open-Source PBX Systems (e.g., Asterisk, FreePBX)**
- **Use of Misconfigured or Permissive VoIP Providers**

The following is a detailed breakdown of each method, presented purely for educational and defensive understanding in cybersecurity investigations:

Method 1

Manual SIP header manipulation via SIP trunking tools

Commercial SIP trunking platforms — such as Twilio, Telnyx, or voip.ms — enable legitimate businesses to control outbound caller ID. However, these platforms can be misused by fraud actors:

01 An account is registered with a SIP trunking service.
02 A local DID (Direct Inward Dialing) number is purchased.
03 SIP credentials (e.g., username, password, SIP server address) are obtained.
04 A softphone (e.g., Zoiper, MicroSIP) is configured with the SIP credentials.
05 A spoofed caller ID is manually entered in the SIP trunk configuration.
06 Outbound calls are made, displaying the falsified caller ID to recipients.

In the following example, a Singapore number (+65 6035 xxx) was used as the DID, and the spoofed caller ID was set to “8144 4444”. When calling an overseas number, the recipient saw “8144 4444” as the caller ID, demonstrating successful spoofing.

Purchasing a DID

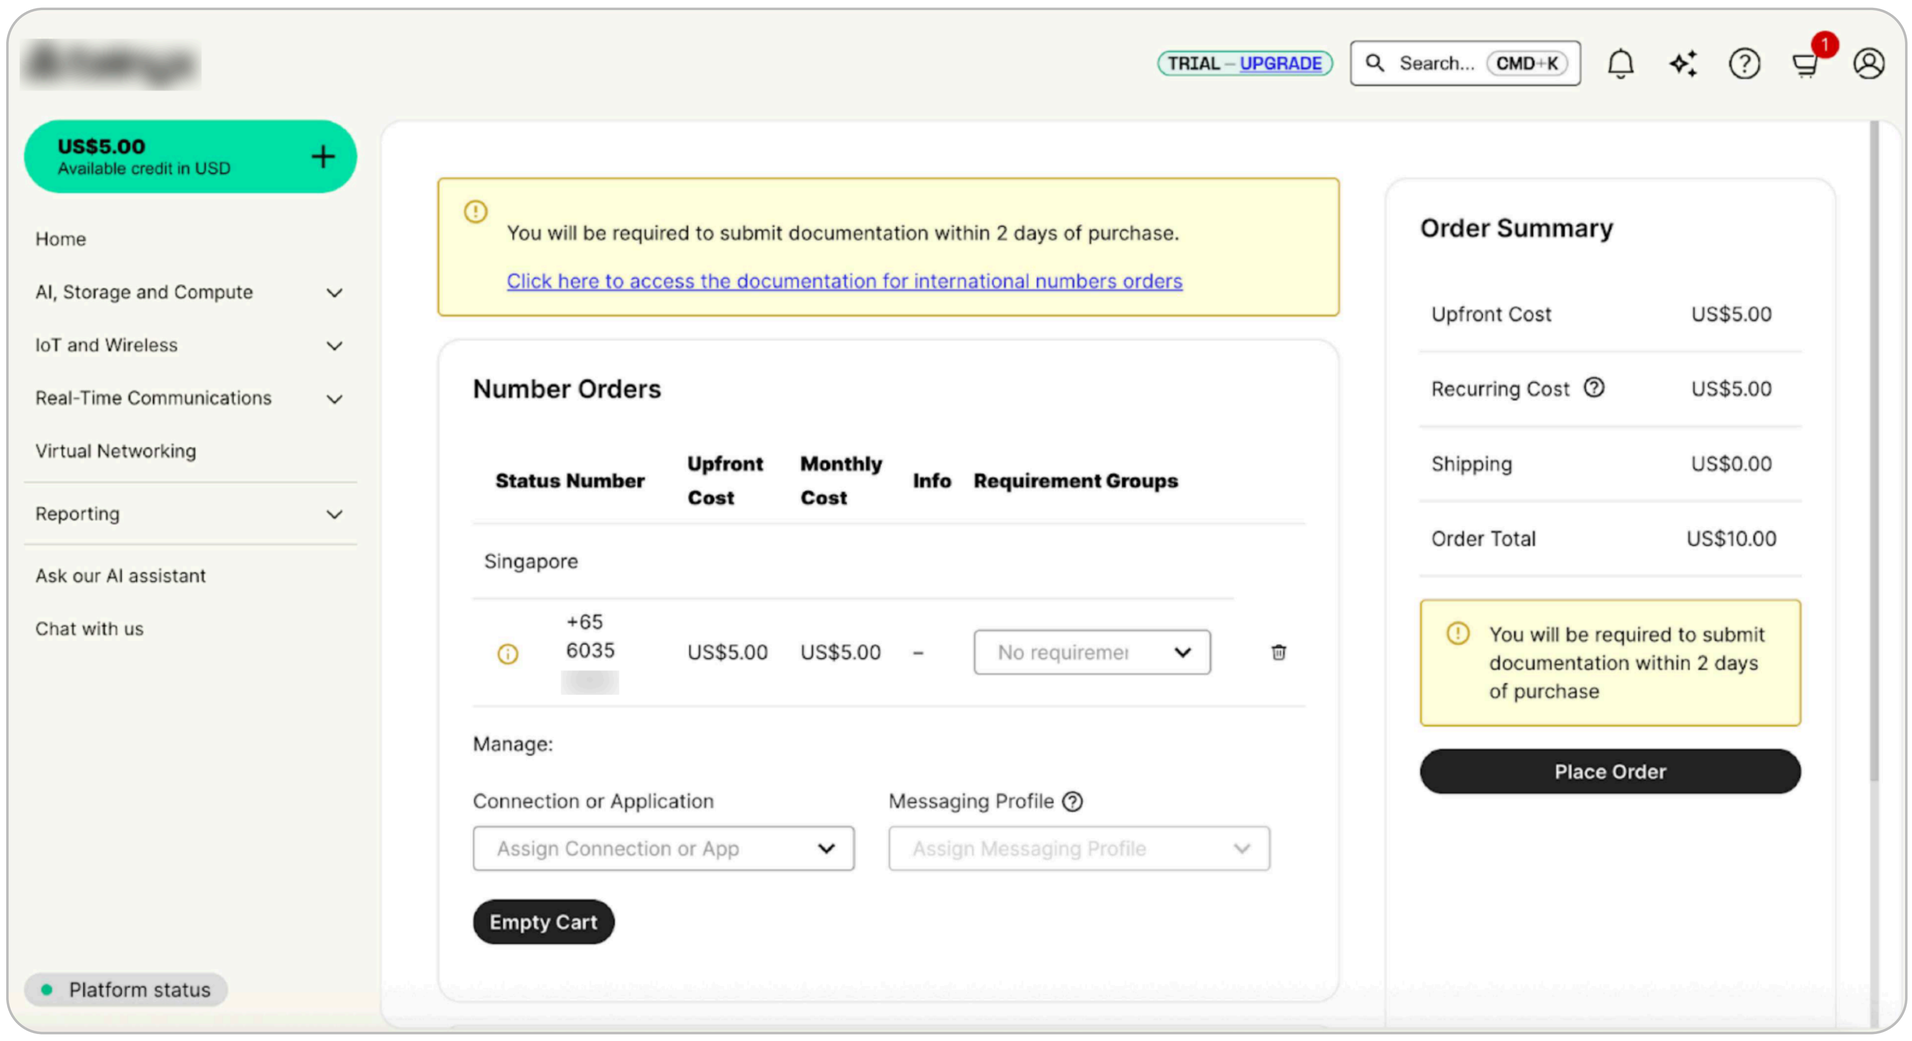

Figure 4.
A screenshot of how a DID can be purchased online.

Get SIP credentials

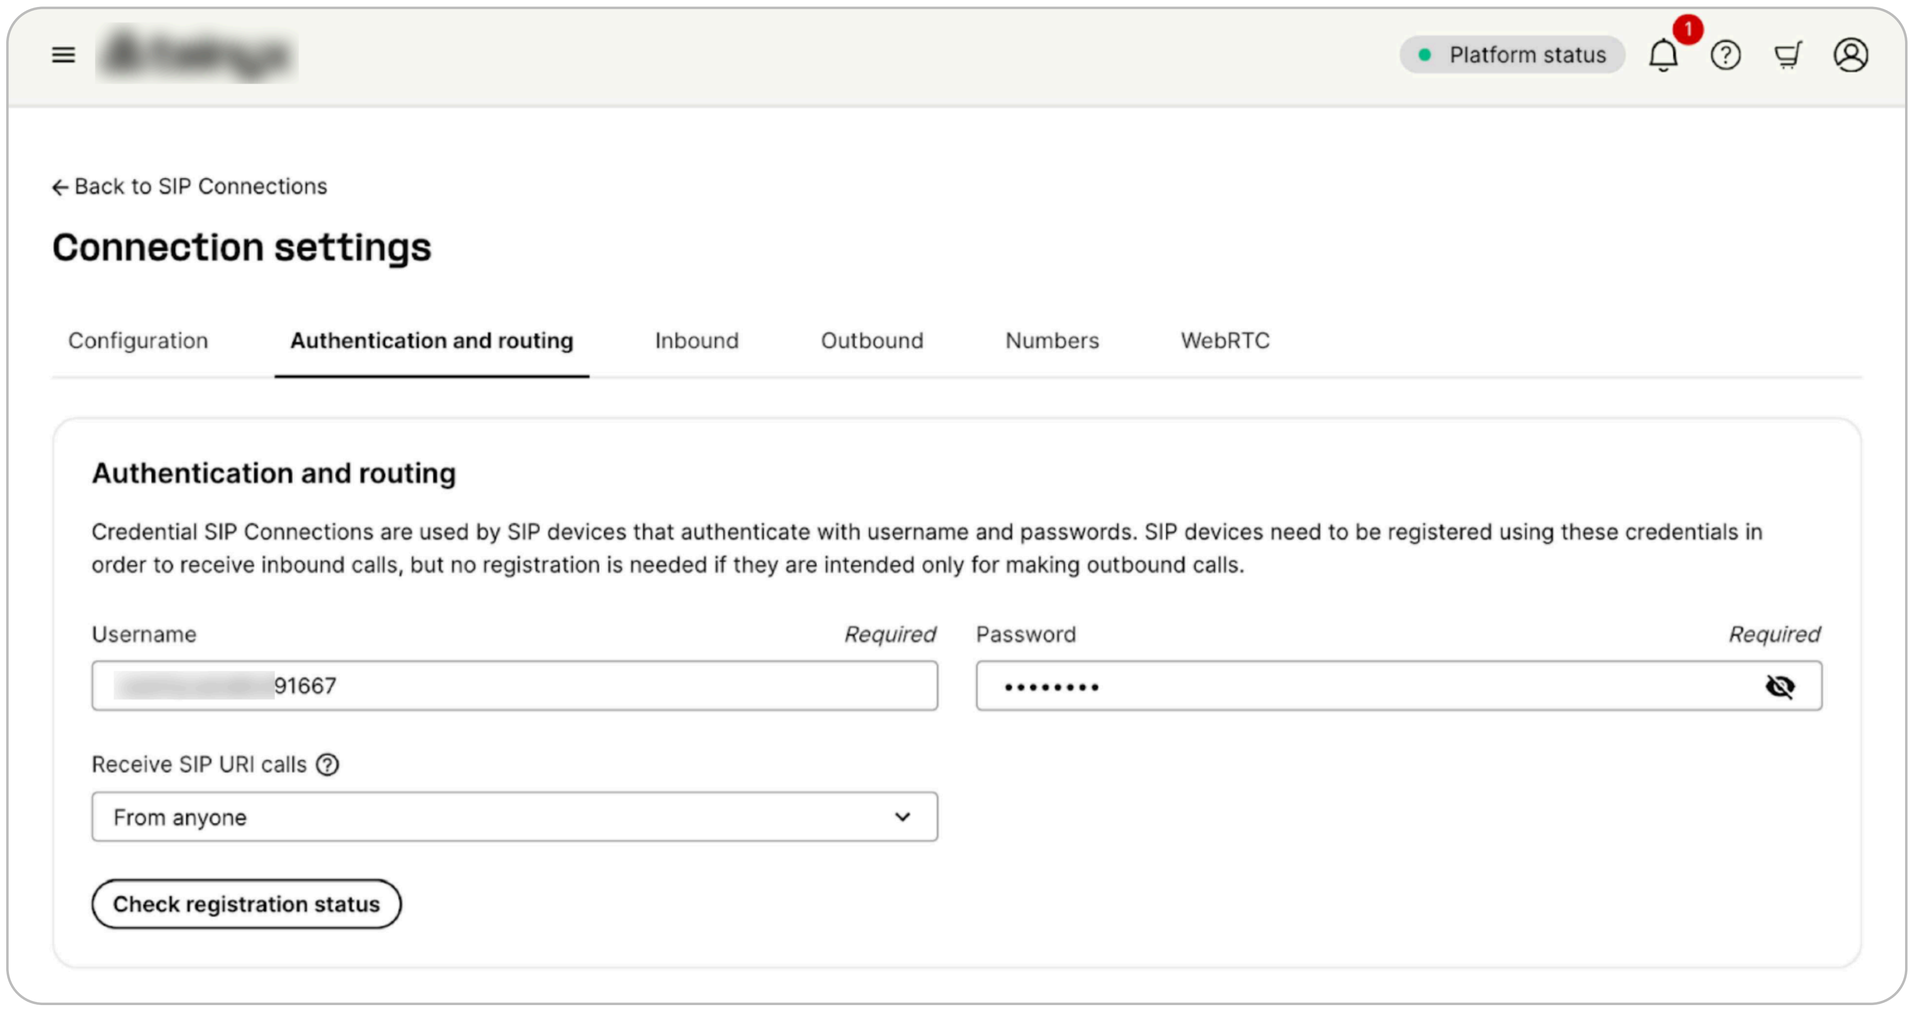

Figure 5.
A screenshot of how the SIP credentials are acquired.

Configure softphone

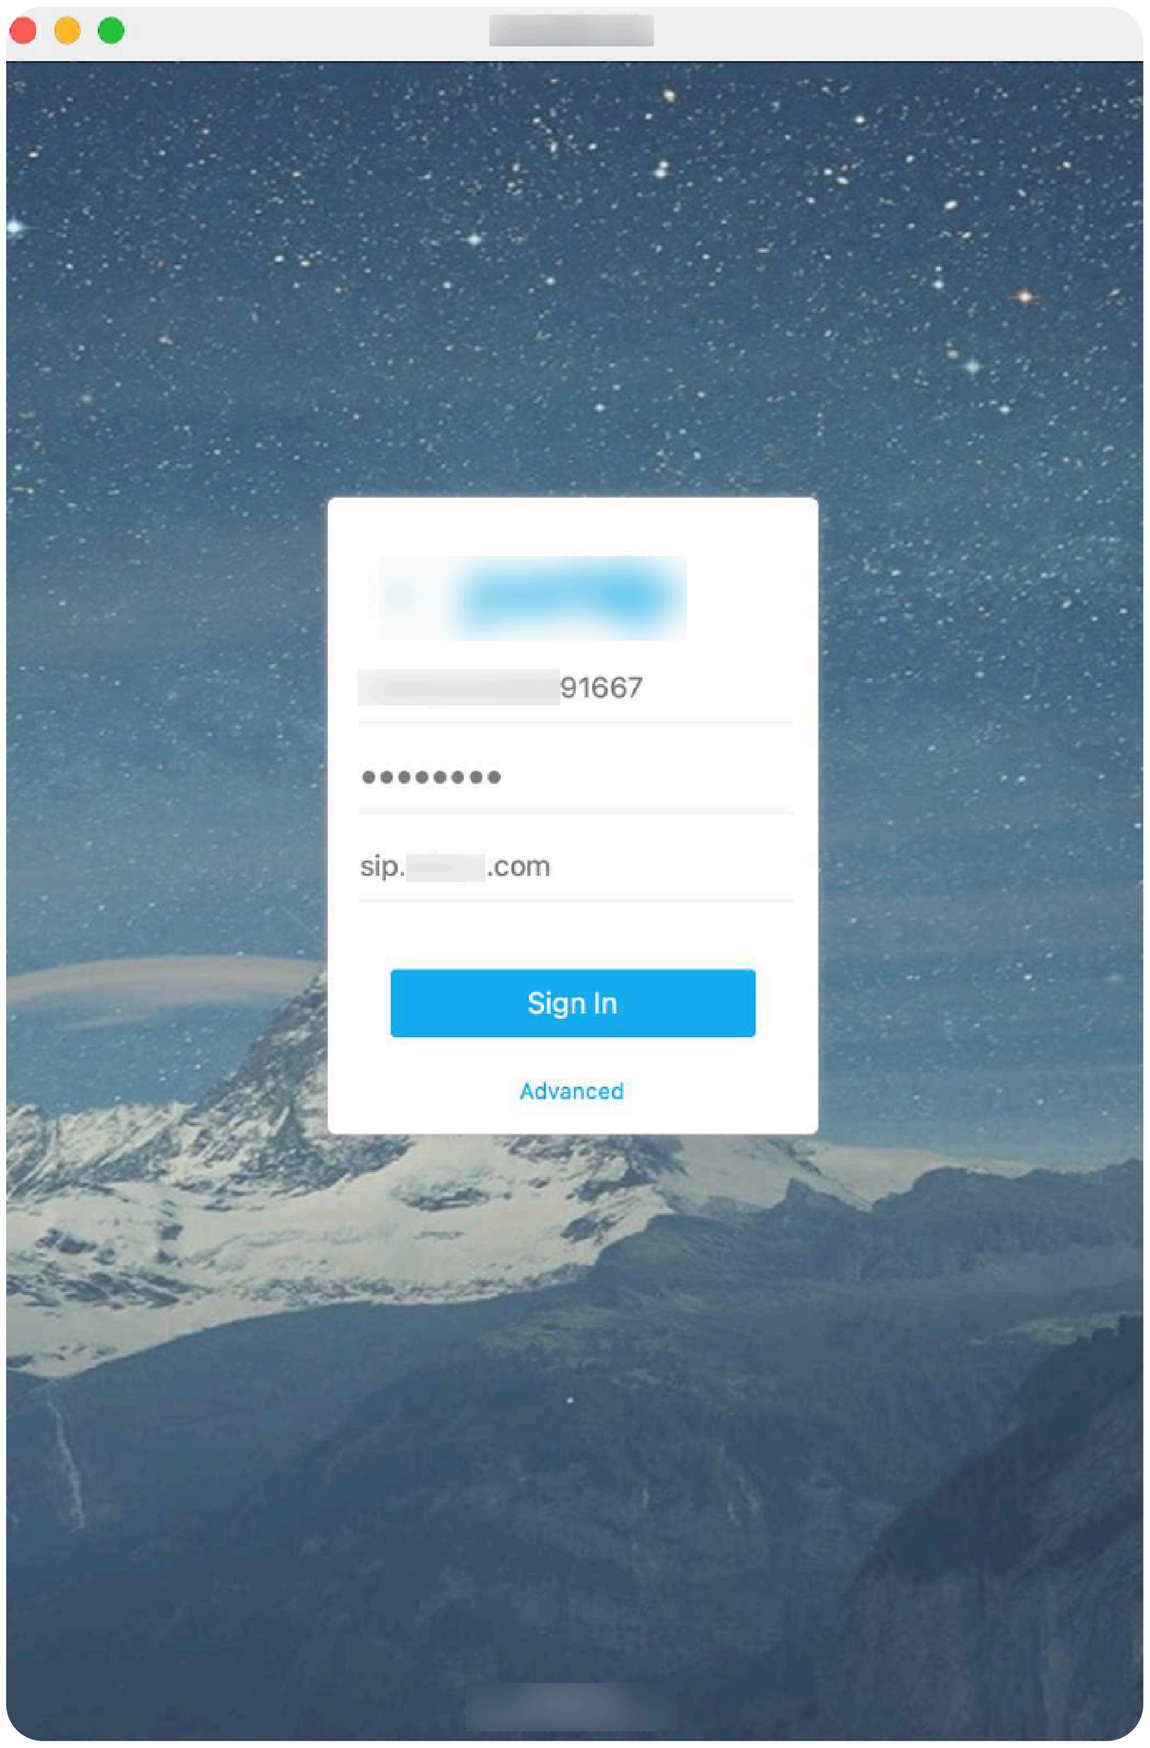

Figure 6.
A screenshot
of a softphone
configuration user
sign-in page.

Custom caller ID

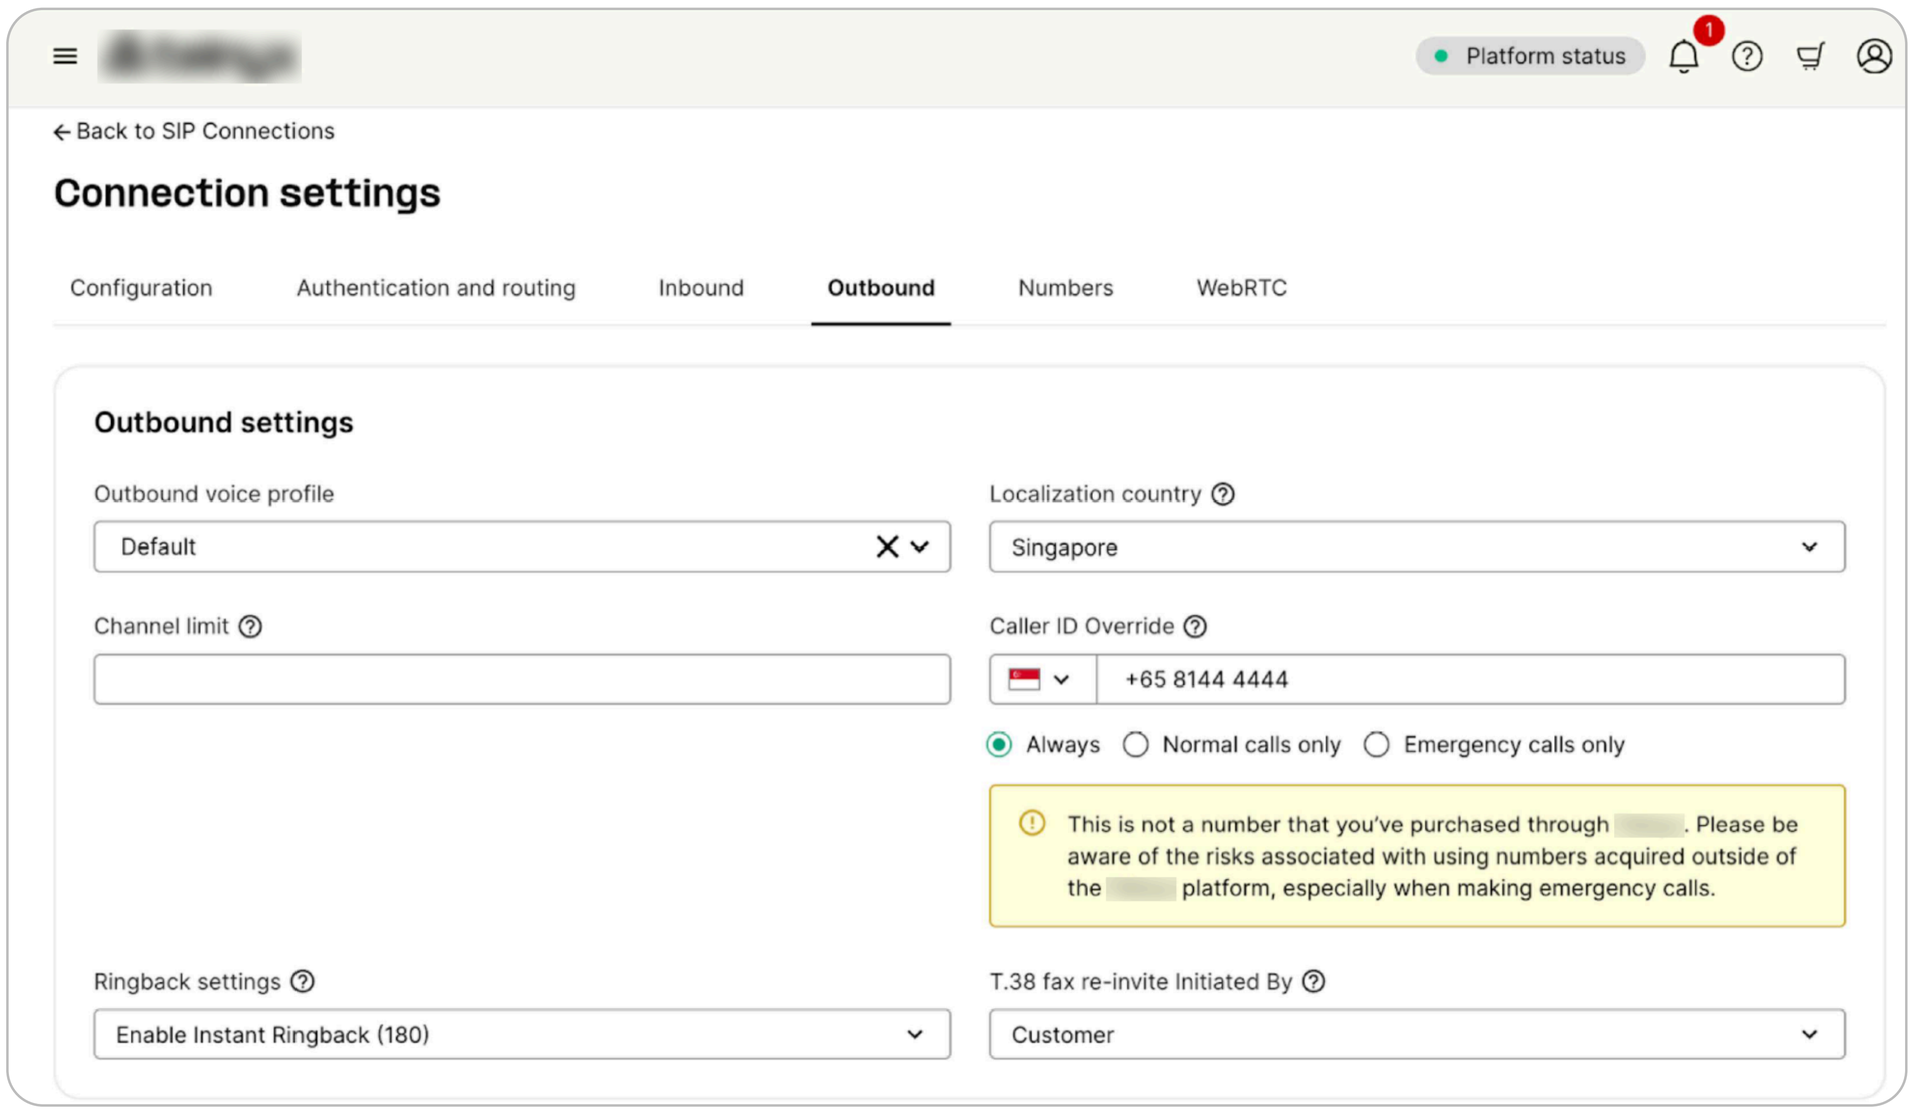

Figure 7.
A screenshot of how the Caller-ID can be overridden.

Calling an overseas number and the spoofed ID received

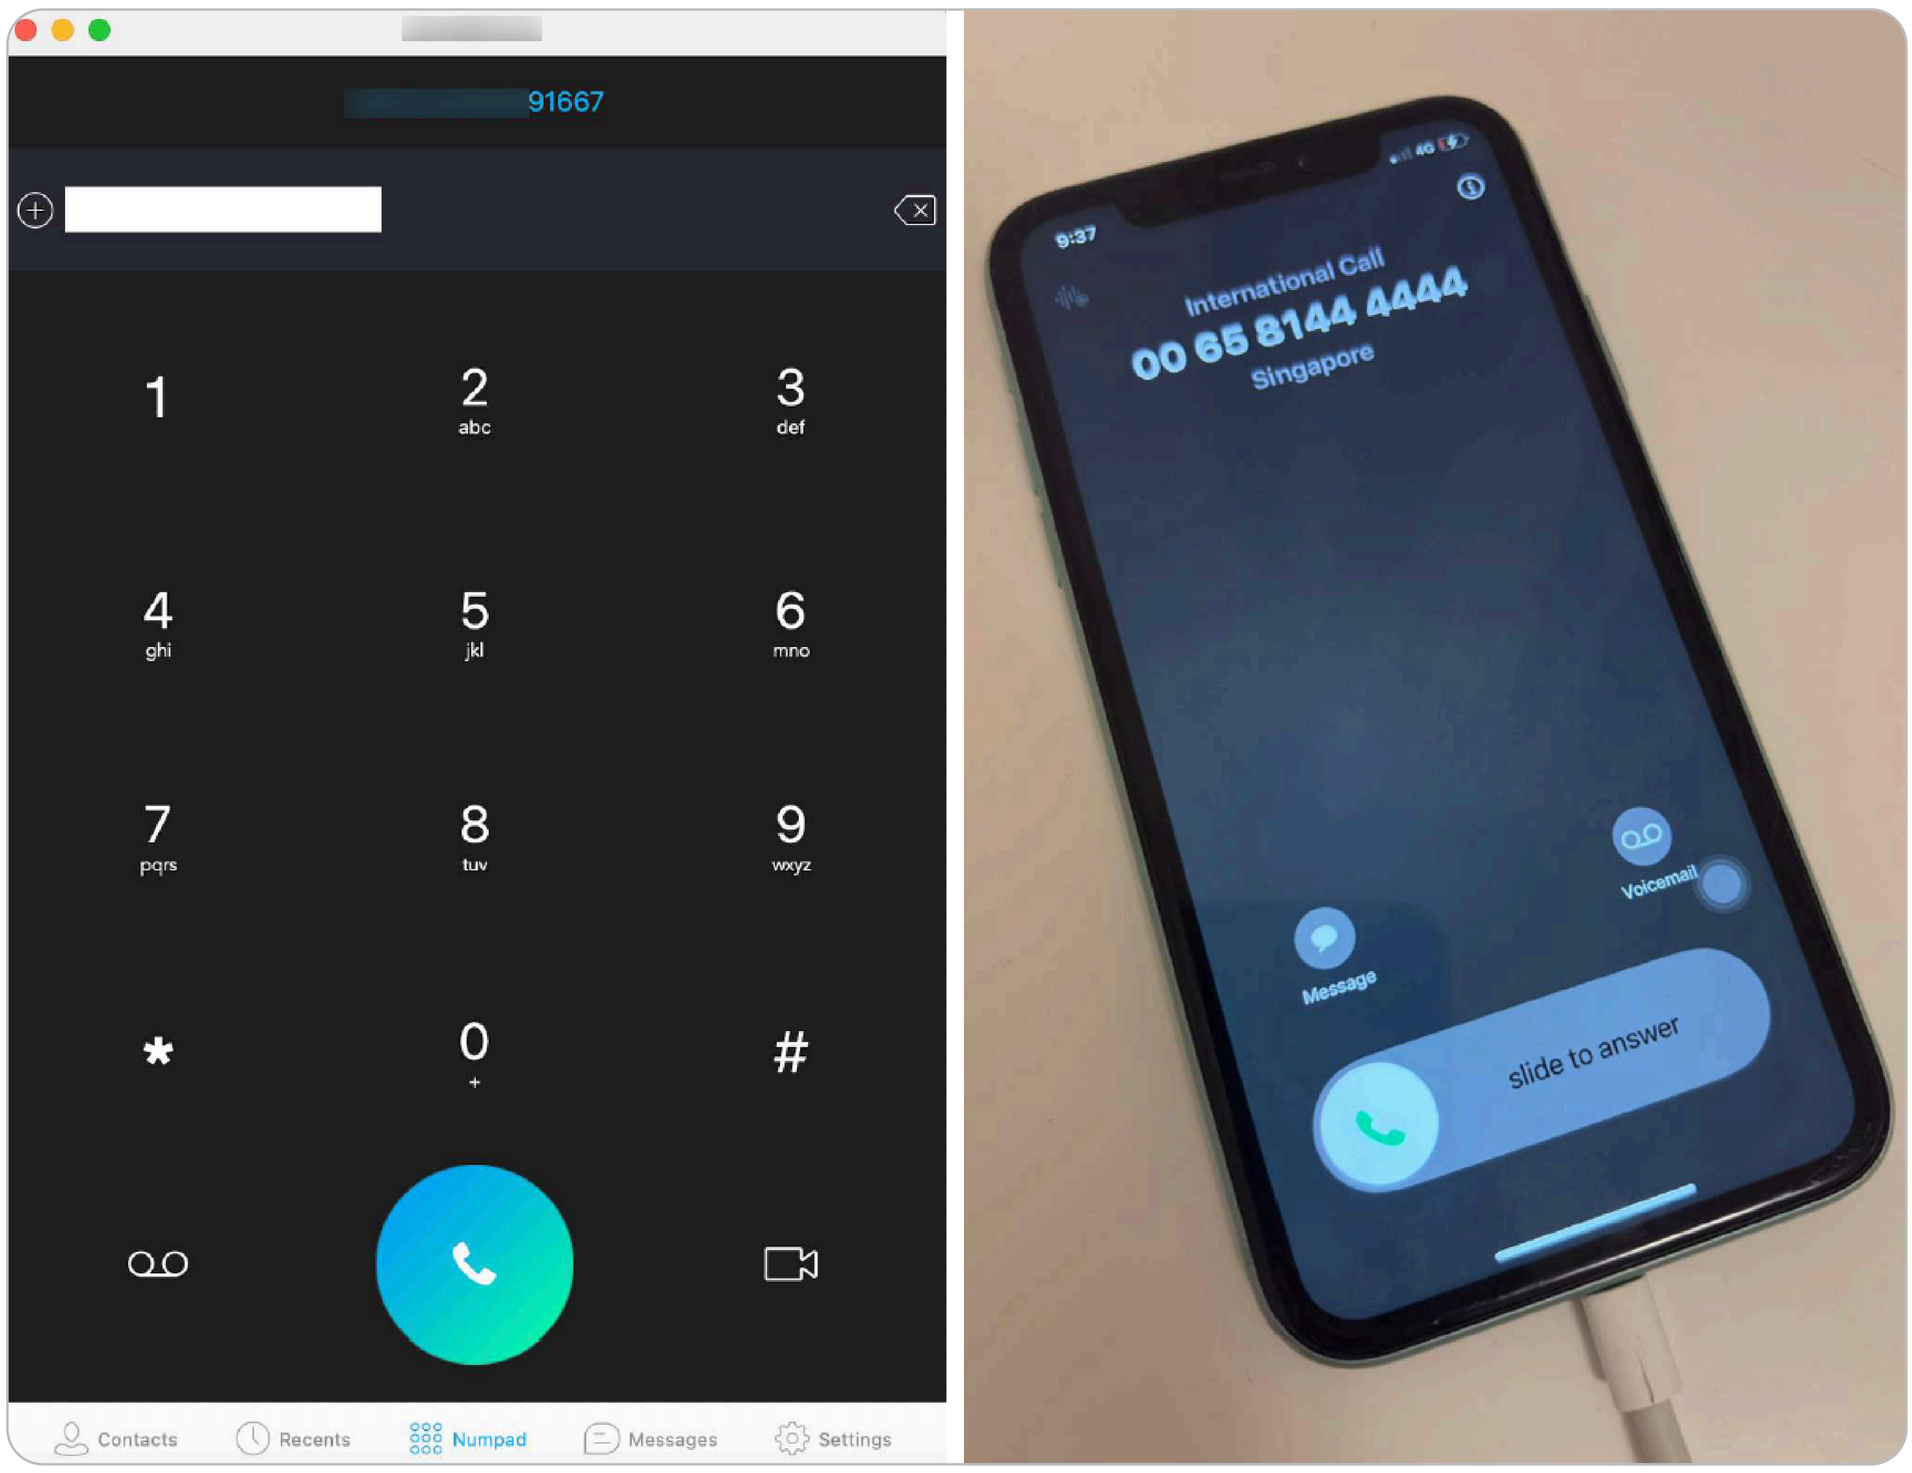

Figure 8. Instead of the DID +65 6035 xxxx, a spoofed number 8144 4444 appears on the recipient's phone screen.

Method 2

Use Open-Source PBX Software (e.g., Asterisk, FreePBX)

This approach requires more technical effort but provides attackers with greater control and stealth, making detection more difficult:

01 A PBX server is deployed on a cloud platform (e.g., OVH, DigitalOcean).
02 One or more SIP trunks are connected to the PBX.
03 Custom extensions are created and outbound caller IDs configured.
04 A local softphone is set up to connect to the PBX.
05 All outbound calls use the configured spoofed caller ID.

Method 3

Use Misconfigured or Permissive VoIP Providers

The above two methods of caller ID spoofing attempts may fail if routed through a VoIP provider that enforces strict caller ID validation —such as requiring number ownership verification. To bypass these controls, fraudsters often leverage **misconfigured or permissive VoIP providers:**

01 The attacker identifies a permissive provider (e.g., outdated SIP gateways, unregulated international carriers).
02 A SIP connection is established with minimal authentication.
03 A spoofed SIP INVITE is transmitted through the provider (via Method 1 or 2).
04 Because no integrity checks are performed, the falsified caller ID reaches the end recipient.

Key Vulnerabilities in the Telecom Ecosystem Enabling Caller ID Spoofing

The spoofing techniques outlined above highlight several systemic vulnerabilities within the telecom infrastructure:

Lack of Number Ownership Verification in SIP Trunking Services

Many SIP services fail to verify whether users have legitimate control over the numbers used in the "From" header, allowing arbitrary numbers to be spoofed

Permissive or Misconfigured VoIP Providers

Weak authentication policies and lack of enforcement of caller ID integrity make certain providers attractive to fraudsters. These systems often accept and forward spoofed INVITE packets without scrutiny.

Absence of End-to-End Caller ID Authentication

Caller ID verification is not consistently enforced across the telecom routing chain. Once a spoofed INVITE is accepted by any upstream provider, the falsified number is often trusted throughout the network — allowing fraudulent calls to appear legitimate to end users.

Organizational blind spots: voice = identity in corporate environments

In many corporate settings — particularly at the executive level — voice-only authorizations remain a prevalent practice. Although precise data on the adoption rate of such procedures is limited, the rising volume of reported deepfake vishing attacks targeting CEOs, CFOs, and other senior personnel indicates that this trust model is still commonly applied across industries.

This reliance on voice as a sole verification factor is often rooted in operational efficiency. High-level decision-makers require speed, and verbal approvals have historically served as a practical shortcut. However, in the context of **AI-generated speech** and **caller-ID spoofing**, this approach introduces a critical vulnerability.

Once a spoofed call reaches an internal staff member — frequently in finance, HR, or executive support — the social engineering process begins. Organizations that permit informal verbal approvals or “exception workflows” for executives are particularly vulnerable. In such scenarios, attackers exploit three layered psychological triggers:

- **“That sounds like my CFO.”** — Voice cloned using AI-based tools.
- **“It came from their number.”** — Caller ID spoofed via VoIP protocols.
- **“They told me not to delay.”** — Urgency introduced to override normal verification procedures.

These factors create a **perfect storm** where:

- Caller ID suggests legitimacy,
- AI-generated voice creates authenticity,
- Urgency disarms the target’s critical thinking.

Under these conditions, voice-based confirmation alone may be sufficient to trigger sensitive actions such as **transaction approvals, privileged access grants, or disclosure of confidential information.** This **blind spot is actively exploited by fraudsters**, and addressing it requires a cultural shift toward **multi-factor verification** even for top-tier executives.

Detection and defense: strategies across stakeholders

Mitigating the threat of AI-powered voice impersonation requires a **multi-layered, cross-sector response** involving telecom providers, corporate institutions, and individual users. The following recommendations are aligned with the systemic vulnerabilities outlined in this report.

Telecom sector: strengthening the infrastructure layer

Implement Number Authentication Protocols	Telecom providers should implement cryptographic call authentication frameworks such as STIR/SHAKEN, which validate that the caller ID matches the originating source of the call. This helps prevent spoofed calls from being accepted as legitimate.
Strengthen Inter-Carrier Trust Models	A trust framework between upstream and downstream carriers should be enforced to validate identity information so that one provider can trust the identity information passed along by another. This prevents spoofed or unauthenticated calls from propagating across networks.
Enforce Caller ID Ownership in VoIP Services	VoIP and SIP trunk providers must verify that customers own the phone numbers they claim to use. Lack of this verification is a core vulnerability that allows attackers to inject calls with arbitrary caller IDs.
Block Unauthenticated or Suspicious SIP Packets at the Ingress Point	Configure infrastructure to reject SIP messages lacking valid authentication headers or originating from suspicious, unauthenticated routes. This limits the surface for spoofed or illegitimate calls reaching end-users.
Integrate Telecom Fraud Protection Systems	Telecom providers should implement AI-driven fraud detection platform capable of monitoring real-time call traffic and identifying anomalous patterns — such as high volumes of short-duration calls from a single source, irregular call routing, or mismatches in usage behavior, and blocking suspicious calls in real time before they reach subscribers.

Corporate environments: zero trust for voice-only approvals

Organizations must evolve their internal processes to reflect the reality that voice only — no matter how authentic-sounding — can no longer be treated as a standalone proof of identity.

Adopt Multi-Factor Verification for Sensitive Actions

High-risk requests — such as fund transfers, access permissions, or credential resets — must require **additional verification steps**, including written approvals, secure tokens, or dual-authorization workflows.

Establish Clear Protocols for Executive Exception Handling

Where executive exceptions are necessary, they should be accompanied by standardized fallback verification procedures. Voice-only approvals should never be accepted without secondary validation.

Integrate Deepfake Vishing Awareness into Staff Training

Cybersecurity training should include modules on AI-generated voice threats, featuring real-world case studies and simulated vishing scenarios to build familiarity with attacker techniques.

Individual users: awareness and personal safeguards

Employees, executives, and individuals can also take steps to recognize and protect against voice impersonation attempts.

Recognize the Clonability of Voice

Individuals should understand that AI can now replicate a person's voice with high fidelity using only seconds of publicly available audio. Familiarity with this reality is essential to resisting voice-based deception.

Understand the Limits of Caller ID

Caller ID should no longer be treated as a reliable indicator of identity. It can be easily manipulated through various spoofing techniques.

Verify Urgent or Unusual Requests Through Secondary Channels

When receiving a call that includes urgent or abnormal requests, users should pause and seek independent confirmation — either by returning the call via a trusted number, using secure messaging applications, or conducting in-person verification where feasible.

Conclusion

In the rapidly evolving threat landscape, deepfake vishing represents a sobering convergence of social engineering and generative AI. By manipulating the most human of identifiers — our voice — fraudsters exploit long-standing trust assumptions embedded in both our technologies and organizational cultures.

As this report illustrates, the illusion of authenticity — amplified by AI-cloned speech and caller ID spoofing — can effectively bypass traditional fraud defenses. High-level executives, often operating under time pressure and informal protocols, present lucrative entry points for attackers. The consequences are not just financial losses, but also reputational damage and operational disruption.

Addressing this threat requires a multi-pronged response. Telecom providers must reinforce infrastructure-level safeguards such as STIR/SHAKEN and SIP verification. Organizations must abandon voice-only trust models in favor of layered, multi-factor verification — especially for sensitive actions. And individuals, regardless of role, must be trained to critically assess unexpected voice requests, even from familiar contacts.

Ultimately, defending against deepfake vishing isn’t only a matter of deploying technical controls — it demands a mindset shift. In this new era, trust must be earned through verification, not assumed through recognition.

GROUP-IB

Group-IB is a creator of cybersecurity technologies to investigate, prevent and fight digital crime.

1,550+
Successful investigations of high-tech crime cases

500+
Employees

60
Countries

$1 bln+
Saved by our client companies through our technologies

#1*
Incident Response Retainer vendor

*According to Cybersecurity Excellence Awards

11
Unique Digital Crime Resistance Centers

Global partnerships

- INTERPOL
- EUROPOL
- AFRIPOL

Recognized by top industry experts

- Forrester®
- Aité Novarica
- kuppingercole ANALYSTS
- Gartner®
- IDC
- Frost & Sullivan

Fight against cybercrime

GROUP-IB.COM
INFO@GROUP-IB.COM

APAC
+65 3159 4398

EU & NA
+31 20 226 90 90

MEA
+971 4568 1785