


Would You Know Your CEO's Real Voice?

Singapore's Growing Deepfake Vulnerability

More than a third can't tell a real message from their CEO apart from an AI-generated one



When the Boss Isn't Really the Boss

AI voice cloning and deepfake technology have made it possible to convincingly imitate a company's most senior voice in minutes, using nothing more than a few seconds of publicly available audio. In Singapore's workplaces, that capability lands on fertile ground: a culture in which instructions from senior leaders carry significant weight, and in which questioning them still comes with a perceived cost.

New research by KnowBe4 surveyed 500 Singapore-based office workers, via Censuswide, between 6–11 August 2026. It reveals a workplace vulnerability that sits squarely between people and technology. More than a third (36%) of respondents say they would not be confident distinguishing a genuine message from their CEO from one generated using AI, such as through voice cloning or deepfake technology. This uncertainty exists at exactly the moment AI-enabled scams are shifting toward impersonating senior leadership, rather than relying on traditional phishing tactics alone.

That uncertainty is compounded by deference. Half (50%) of respondents say it would feel uncomfortable to challenge a request from a senior executive, and 52% say they would worry about negative consequences if they questioned a senior leader. A hierarchy that discourages pushback, combined with a growing wave of convincing AI impersonation, leaves Singapore's businesses with a structural weak point that technology alone can't close.

This vulnerability has two distinct faces. At the bottom of the org chart, fear of challenging a senior leader keeps people quiet even when something feels wrong. But perhaps most concerning of all is what's happening at the top: the employees who are most confident in their ability to recognise a fake, and who have received the most training to do so, are not the ones proving most resistant to it in practice. This report unpacks both halves of that problem: why deference to authority is becoming a vector for AI-enabled fraud, and where training is and isn't reaching the people who need it most.

What's Inside

- > **03 Key Findings**
- > **04 Fear at the Bottom: A Culture Where Senior Leaders Aren't Questioned**
- > **05 Fear at the Bottom: Weighing the Wrong Risk**
- > **06 Overconfidence at the Top: The Message That Looks Real Enough to Act On**
- > **07 Overconfidence at the Top: Confidence vs Reality**
- > **08 Training Gaps and the Case for Workforce Security**
- > **09 Building Trust That Can't Be Faked**

More than a third of Singapore's office workers say they would not be confident distinguishing a genuine message from their CEO from one generated using AI

Key Findings



36%

of office workers are not confident they could distinguish a genuine message from their CEO from one generated using AI



50%

Half say it would feel uncomfortable to challenge a request from a senior executive



52%

say they would worry about negative consequences if they questioned a senior leader



45%

would be equally worried about delaying a request from a senior leader as they would about accidentally causing a cybersecurity incident — and 9% would rather risk the breach than delay the boss



78%

Senior managers are the most confident group at spotting a fake (78%) but are also the most likely to comply with a cloned CEO voice before verifying it (35%)

Only


26%

have received training specifically on AI-generated scams involving senior leaders

“In Singapore, hierarchy works like a trust shortcut, one that AI-enabled scams are now designed to exploit. A cloned voice or a convincing message from ‘the boss’ doesn’t need to fool the technology. It only needs to activate an instinct not to question authority. Closing that gap requires training that changes behaviour at every level of seniority, not just awareness at the top.”

Dr. Kawin Boonyapredde,
CISO advisor at KnowBe4

Fear at the Bottom: A Culture Where Senior Leaders Aren't Questioned




Attitudes toward hierarchy in Singapore's workplaces are split almost down the middle. Nearly half (48%) of respondents agree that employees are expected to follow instructions from senior leaders without questioning them, and half (50%) agree it would feel uncomfortable to challenge a request from a senior executive. Over half (52%) agree they would worry about negative consequences if they questioned a senior leader. At the same time, almost half (48%) agree employees can safely question senior leaders without fear of repercussions, and the same share (48%) agree their organisation encourages employees to question unusual requests, regardless of who they come from.

Respondents from Gen X (born 1965–1980) are more likely than those from Gen Y, also known as Millennials (born 1981–1996), and Gen Z (born 1997–2012) to strongly agree that employees are expected to follow instructions from senior leaders without questioning them (15% vs 13% and 8% respectively), and men are more likely than women to agree with this overall (52% vs 44%).


When asked what would make them hesitate before questioning a request from senior leadership, the top reasons respondents give are:

What Makes Singapore's Office Workers Hesitate to Question Senior Leadership



A further 17% cite a belief that senior leaders know best, while 14% say they wouldn't hesitate at all. Gen Z (46%) and Millennials (40%) are more likely than Gen X (33%) to say fear of delaying important work would make them hesitate, and respondents at multinational companies are more likely than those at local or national companies to cite fear of appearing disrespectful (47% vs 36%).

Fear at the Bottom: Weighing the Wrong Risk




Over two in five (45%) office workers surveyed say they would be equally concerned about delaying an important request from a senior leader and accidentally causing a cybersecurity incident. A further 9% say they would be more concerned about delaying a request from a senior leader than about causing a cybersecurity incident. Just 41% say they would be more concerned about causing a cybersecurity incident, meaning the majority of Singapore's office workers do not treat a security incident as the greater risk.

Gen X (53%) are more likely than Millennials (42%) or Gen Z (40%) to say they are equally worried about both, and those who are fully office-based are more likely to say this than those in a hybrid working arrangement (50% vs 40%).

Asked directly what creates the greatest cybersecurity risk in the workplace today, respondents point most often to gaps in awareness and process rather than technology itself:


The Greatest Cybersecurity Risks in Singapore's Workplaces Today



Millennials and Gen X are most likely to say the greatest risk is employees not recognising AI-generated scams (28% and 24% respectively), while Gen Z is most likely to point to a lack of cybersecurity training (26%). Junior managers (34%) are more likely than middle managers (24%) and senior managers (21%) to say the greatest risk is employees not recognising AI-generated scams, and fully office-based respondents are more likely than those in a hybrid arrangement to point to a lack of cybersecurity training (22% vs 14%).


On paper, most workplaces say they support employees pushing back on suspicious requests. Over three in five (62%) office workers agree their organisation empowers employees to question suspicious requests from senior leaders, even if they appear urgent, while 12% disagree. Again, respondents at multinational companies are more likely to agree than those at local or national companies (67% vs 59%), and senior managers are far more likely to agree than middle managers, junior managers and clerical workers (82% vs 63%, 50% and 51% respectively). That gap, between what organisations say on paper and what employees actually feel able to do, is the first problem. The second sits further up the chain, where confidence outpaces caution.

Overconfidence at the Top: The Message That Looks Real Enough to Act On



Respondents were asked to imagine an urgent WhatsApp message that appeared to come from their CEO. The message asked them to transfer confidential company information immediately, since the CEO was travelling and unable to access company systems. Most say they would take a verification step first — but not all.

What Singapore's Office Workers Would Do First If They Received a Suspicious CEO WhatsApp Message




Gen X respondents are the most likely to say they would send the information immediately, and Gen Z the least likely (10% vs 6%). Respondents at multinational companies are more likely than those at local or national companies to say they would report it to IT/security (28% vs 23%), as are those in hybrid working arrangements compared with those fully office-based (29% vs 21%).

A second scenario tested how respondents would react to a voice that sounded exactly like their CEO's, asking them to urgently bypass a normal approval process because of an important business issue. Just over three in five (61%) say they would be unlikely to comply before independently verifying the request, assuming the voice sounded completely authentic. However, an alarming 23% say they would likely bypass normal approval processes, including 5% who say they would be very likely to do this. Millennials are the most likely to say they would comply (26%), and Gen X the least likely (20%); men are also more likely than women to say this (26% vs 20%).

Despite being more likely to have received training specifically on AI-generated scams involving senior leaders, senior managers (35%) are more likely than middle managers (23%), junior managers (21%) and clerical workers (16%) to say they would comply with a cloned CEO voice before independently verifying the request.

Overconfidence at the Top: Confidence vs Reality




Just under three in five (59%) office workers say they would be confident telling a genuine message from their CEO apart from an AI-generated one, while over a third (36%) would not. Men are more likely than women to say they would be confident (63% vs 56%), and Gen Z (65%) are more likely than Millennials (59%) and Gen X (59%) to say the same.

Confidence is highest among senior managers. Senior managers (78%) are far more likely than middle managers (61%), junior managers (48%) and clerical workers (49%) to say they would be confident distinguishing a real message from an AI-generated one. Yet as Section 3 shows, that same group, senior managers, is also the most likely to say they would comply with a cloned voice before verifying it.

Confidence vs Compliance, by Seniority


Data (Confident distinguishing AI from genuine / Would comply with cloned CEO voice before verifying):



The data points to an inverse relationship between seniority and caution. As confidence in spotting a fake rises with rank, so does the likelihood of acting on one before checking whether it's real, an unusual pattern in an era when the most senior voices in a business are also the ones most likely to be impersonated.

Training Gaps and the Case for Workforce Security

Cybersecurity training in Singapore's workplaces remains inconsistent, and where it exists, it is not reaching everyone equally. Just over a quarter (26%) of respondents receive training specifically on AI-generated scams involving senior leaders, while 27% receive general phishing training only, 17% have received training but don't recall whether it covered AI scams, 11% say they rarely receive cybersecurity training, and 12% receive no cybersecurity training at all.



Respondents at multinational companies are more likely than those at local or national companies to receive training specifically on AI-generated scams involving senior leaders (34% vs 20%). That training is also concentrated at the top of the org chart: senior managers (39%) are far more likely to have received it than middle managers (23%), junior managers (22%) or clerical workers (12%), the same tier that Section 3 shows is most likely to comply with a cloned voice before verifying it.

Separately, when asked which AI-enabled threats concern them most, respondents rank AI-generated phishing emails or impersonation messages highest (61%), followed by AI-generated documents or fraudulent document-sharing requests (56%), deepfake video calls (55%) and voice cloning (53%). Gen Z stands out as the generation most concerned specifically about voice cloning (66%), the exact scenario in which seniority and confidence are shown to increase, not reduce, risk.

Taken together, these findings describe two distinct problems, not one. Fear at the bottom keeps junior and middle-tier staff from challenging a request that feels off, even when their organisation says it welcomes pushback. Overconfidence at the top means training and self-assurance are concentrated where compliance risk is highest. Awareness alone is not closing either gap. What's needed is a modern approach to workforce security: one that measures how people actually behave under the specific pressure of an urgent request from someone senior, and coaches that behaviour continuously, at every level of seniority, rather than treating training as a one-time box to check.



Building Trust That Can't Be Faked


Singapore's office workers are not oblivious to the threat AI-enabled impersonation poses. But this research points to two separate vulnerabilities, not one. Further down the org chart, fear of challenging a senior leader keeps people silent even when they sense something is wrong. Further up, confidence and training are concentrated among the group most likely to act on a convincing fake before checking it's real. Both are exactly what AI-enabled impersonation scams are built to exploit.

For Singapore's organisations, the message is clear: these two problems need two different fixes. Junior and middle-tier staff need genuine psychological safety to question a senior-sounding request without fear of consequence. Senior leaders need better verification habits, not more confidence, since the data shows confidence alone doesn't reduce risk. Until questioning an urgent request from the top becomes as automatic as double-checking a suspicious link, at every level of seniority, Singapore's businesses will keep failing the deepfake test.

About KnowBe4

KnowBe4 empowers the modern workforce to make smarter security decisions every day. Trusted by more than 70,000 organizations worldwide, KnowBe4 is the pioneer of digital workforce security, securing both AI agents and humans. The KnowBe4 Platform provides attack simulation and training, email & collaboration security, and agent security powered by AIDA (Artificial Intelligence Defense Agents) and a proprietary Risk Score. The platform leverages 16-years of behavioral data to combat advanced threats including social engineering, prompt injection, and shadow AI. By securing humans and agents, KnowBe4 leads the industry in workforce trust and defense.

More information at www.KnowBe4.com.



Methodology

This report is based on an independent online survey conducted by Censuswide among a sample of 500 Singapore-based office workers who use email/computers for work. The data was collected between 06.08.2026–11.08.2026. Censuswide is a member of the Market Research Society (MRS) and the British Polling Council (BPC), and a signatory of the Global Data Quality Pledge. We adhere to the MRS Code of Conduct and ESOMAR principles.

Generational classifications used throughout this report are as follows: Generation Z (born 1997–2012), Millennials (born 1981–1996), Generation X (born 1965–1980).



KnowBe4, Inc. | 33 N Garden Ave, Suite 1200, Clearwater, FL 33755
855-KNOWBE4 (566-9234) | www.KnowBe4.com | Sales@KnowBe4.com

Other product and company names mentioned herein may be trademarks and/or registered trademarks of their respective companies.