The security value of voice biometrics

How voice biometrics can reduce fraud
for IVRs and contact centers.

NUANCE

Table of contents

Abstract

Nuance voice biometric solutions have consistently reduced exposure to fraud originating in the IVR or contact center. This ability to reduce fraud has been predominantly revealed within financial institution deployments of voice biometrics, where the phone customer service channel represents a relatively easy target to conduct account takeovers and perform fraudulent fund transfers. Financial institutions that have deployed Nuance voice biometric solutions according to best practices have seen a dramatic reduction in fraud, in some cases up to 10 fold.

This white paper outlines how Nuance voice biometric solutions eliminate several security vulnerabilities that exist in IVRs and contact centers. These include the weaknesses associated to PIN credentials used to secure IVR self-serve functions, as well as the vulnerabilities inherent in an agent security question verification process. This paper also provides an overview of how a fraudster may attempt to compromise a voice biometric system, while listing the counter-measure Nuance voice biometric solutions offer to mitigate such threats.

It is important to keep in mind that no single security solution can prevent all malicious attacks. However, it is clear that migrating from PINs and security questions to voice biometrics in order to validate caller identity can significantly reduce fraud within the IVR and contact center environments.

How do fraudsters typically compromise security in IVRs today?

Most IVRs use a knowledge authentication factor to secure caller access, the PIN being the credential most often used. The methods listed below to compromise a PIN also apply to other knowledge factors.

Brute force attack

The four digit PIN is one of the weakest security credentials, due to the ease for a malicious user to compromise a system protected by a PIN without the need to possess any technical knowledge, or any knowledge of the legitimate account holder. PIN protected IVRs are often the targets or organize fraud groups that can rapidly compromise large numbers of accounts with a small number of calls to the IVR. The problem is revealed by a 2012 study conducted by DataGenetics on the vulnerability of PINs.[1] The study shows that 10.7% of four digit PINs are "1234". This means that a fraudster calling into an IVR, accessing random accounts would need to conduct an average of 10 attempts to compromise an account! The malicious individual then has full access to the account, and can repeatedly commit fraud without being detected.

Although organizations can block the most commonly used PINs, the study also reveals that beyond sequential numbers and repeating numbers, people tend to select PINs where the numbers form patterns on keyboard, or where the number represents a date that is significant to the caller.

This PIN selection behavior by legitimate account holders render brute force attacks quite effective, and not surprisingly have limited organizations in the self-serve functions that can be offered in a system protected uniquely by a PIN. In effect, organizations typically require a caller to perform higher-risk transactions through an agent, who will ask a series of security questions. As is demonstrated in the next section of this paper, the agent conducted security question process is the weakest form of security, even more vulnerable than the PIN, which explains why the bulk of fraud is conducted via the agent.

Compromising the PIN database

A PIN, like any other knowledge factor used for authentication to an IVR, is stored in a database. If the database is compromised, a malicious user has unlimited access to accounts. Although properly designed systems have numerous security measures in place to prevent a security breach, there are numerous documented cases of breaches occurring. Some cases involve ingenious hackers finding ways to bypass the security measures. Other cases involve employee error, for example an erroneous transfer of PIN credentials through e-mail. No matter how the PINs are compromised, once in the hands of a malicious individual, the potential for large scale financial losses are enormous.

Most Commonly Used PINs

	PIN	FREQ
#1	1234	10.713%
#2	1111	6.016%
#3	0000	1.881%
#4	1212	1.197%
#5	7777	0.745%
#6	1004	0.616%
#7	2000	0.613%
#8	4444	0.526%
#9	2222	0.516%
#10	6969	0.512%
#11	9999	0.451%
#12	3333	0.419%
#13	5555	0.395%
#14	6666	0.391%
#15	1122	0.366%
#16	1313	0.304%
#17	8888	0.303%
#18	4321	0.293%
#19	2001	0.290%
#20	1010	0.285%

1 http://www.guardian.co.uk/money/blog/2012/sep/28/debit-cards-currentaccounts
2 Study conducted by University of Cambridge on behalf of UK's Consumer Association, 2010. 1,000 participants were surveyed online.

Credential sharing

Any knowledge factor can be shared, and PINs are no exception. Banking card PINs tend to be widely shared, as legitimate account holders allow a spouse, child, relative, caregiver or colleague to perform transactions with the banking card. In fact, a survey conducted in 2010 by the University of Cambridge revealed that 10% of ATM card users admitted to sharing their PIN.[2] PINs used in IVRs suffer the same fate, and this makes accounts vulnerable should someone with knowledge of the PIN have malicious intent. PIN sharing is a commonly acknowledged problem, but prior to voice biometrics, there were no practical alternative to enhance security within the IVR context. Since the PIN doesn't identify the individual in any way, once dialed in the IVR, the malicious user has complete access to the legitimate customer's account. As such, PIN sharing provides the malicious user with a 100% success rate of conducting an attack.

Phishing

Phishing is an every-increasing technique that malicious individuals undertake to compromise credentials, such as PINs. A study by Gartner indicates that a random Phishing attack yields a 3% data collection success rate, meaning that if 100 e-mails are sent to collect PINs, a hacker will on average collect 3 valid PINs.[3] However, if the malicious individual conducts a context-aware phishing attack, the success rate can reach 72%, as demonstrated by a study conducted by Indiana University.[4] As such, Phishing attacks are one of the preferred choices by malicious individuals to compromise systems that are protected by PINs and Passwords.

Vhishing

The art of convincing people to divulge confidential information is a technique that is leveraged by the more sophisticated fraudster. It requires calling a legitimate account holder and deceiving them in order to collect their credential, such as their PIN. This is a high-effort activity and tends to yield a lower success rate when compared to a simple brute force attack or phishing (via e-mail). However, if the fraudster is intent on compromising a specific individual's account, Vhishing is often the fallback technique when alternative methods are unsuccessful.

PIN reset function

PINs have a tendency of being forgotten. In the IVR, this is particularly a problem as call frequency tends to be low, with the exception of callers that use telephone banking to conducting their regular transactions. According to data collected from Nuance customers, the typical PIN authentication failure rate ranges between 30% and 70% of all PIN authentication attempts. The bulk of these failures are caused by legitimate callers forgetting their PIN. This is why a process to reset PINs is readily available to callers. In most cases, the caller is transfered to an agent, and is asked to answer a series of security questions in order to reset their PIN. The PIN reset function can be easily leveraged by malicious users to reset an account holder's PIN and then use the new PIN to commit fraud. The vulnerabilities associated to the agent handled security questions process are reviewed in the next section.

30%-70%

Typical PIN authentication failure rate ranges between 30%-70% of all PIN authentication attempts.

3 Gartner Inc. Gartner study finds significant increase in e-mail phishing attacks, April 2004.
4 http://www.indiana.edu/~phishing/social-network-experiment/phishing-preprint.pdf

How do fraudsters typically compromise security in contact centers today?

Calls centers typically use a series of knowledge questions to verify a caller's identity. Agents will ask for information such as the caller's address, phone number, mother's maiden name, social security numbers (SSN), or last transaction amount and vendor. If the caller answers the questions correctly, the agent considers the caller's identity validated and any transactions can then take place. The vulnerabilities listed below showcase some of the methods used by fraudsters to overcome this security process.

Internet search

Many of the answers to the security questions asked by contact center agents can be easily found on the internet. A moderately sophisticated hacker can find the answers to the majority of security questions by accessing social media sites, such as Facebook and LinkedIn, often collecting phone number, date of birth, and complete employment history. With this data, a malicious user can then use algorithms to deduce an individual's SSN.[5] Furthermore, collecting basic information about an individual online makes the task of guessing answers to security questions easy, as was shown by a study at Carnegie Mellon University in 2009. It demonstrates that typically used security questions are vulnerable, in some cases they can be guessed with a 48% accuracy.[6] If any of the information being requested by a contact center agent is too difficult to find online or to guess, a malicious user can always purchase the information from an information broker.[7] Credit card transactions statements and SSN can be purchased from these brokers, some who claim 99% data accuracy, and the ability to provide data for any individual for as little as $1 per data point requested.

Hacking

Conducting searches on the web and purchasing data from information brokers are effective methods at compromising a specific individual's account. However, the sophisticated malicious user will attempt to hack an entire database containing all of the answers to security questions in order to conduct large scale fraud. High-profile cases of such hacking cases abound, one of the most recent US cases involves a compromised Department of Revenue server in South Carolina that exposed 3.6m SSN numbers, as well as other information contained in tax returns including credit card number, debit card number, address, names and date of births.[8] With such information, malicious users can then reach a contact center, provide all of the appropriate answers to security questions, and perform a complete account takeover for a larger number of account holders. Within hours, an organized fraud gang can defraud hundreds of accounts, providing fraud prevention specialists little time to react.

5 http://arstechnica.com/science/2009/07/
   social-insecurity-numbers-open-to-hacking/
6 http://research.microsoft.com/pubs/79594/
   oakland09.pdf
7 http://archive.newsmax.com/archives/
   articles/2005/4/4/155759.shtml
8 http://www.thestate.com/2012/10/26/
   2496396/south-carolina-taxpayers-privacy.
   html#.URlsjaXC1qw

67%

of social engineering attempts at bank contact centers were successful.

Social engineering
Customer service agents are instructed to minimize AHT and deliver a positive customer experience to callers. Fraud prevention is for most agents not their number one priority. This creates an opportunity for fraudsters to compromise an account through social engineering. According to a study conducted by Global Reviews in 2011, about 67% of social engineering attempts at bank contact centers were successful.[9] Contact center agents lack the training and incentives to detect social engineering attempts. Organizations that have required agents to comply with stringent security procedures have seen disastrous impacts on customer care. Indeed, one financial institution reported that over 20% of legitimate callers were unable to receive service. As such, organizations tend to require a minimum of inconvenience to the caller in order to reduce the impact of the security process on the customer experience. However, this creates an important security vulnerability that fraudsters leverage at an ever increasing frequency.

Why voice biometrics provides a superior security solution over PINs or agent security questions
As presented in the previous sections, PINs and security questions have inherent security vulnerabilities when used in IVR and contact center environments. This section reviews why voice biometrics offers improved security over these identity verification methods.

A compromised voiceprint is unusable for account access
A voiceprint is a hashed string of numbers and characters that represent how a specific individual's voice rates on the myriad of characteristics being measured. As such, a compromised voiceprint has no value to a hacker. It cannot be used to authenticate to a system. Neither can it be used to reverse engineer someone's voice. This inherent characteristic of voice biometrics provides a fundamental security benefit over any knowledge based authentication method. Compromised credential databases pose a massive risk to security as such breaches have the potential of inflicting the largest financial losses. Voice biometrics is not vulnerable to such a large scale attack. This limits fraudsters to attempting to compromise individual accounts. This security benefit of voice biometrics dramatically reduces the fraud risk that an organization faces within the IVR and contact center service channels.

Proactive detection of known fraudsters
Each time a fraudster speaks within an IVR or to a contact center agent, the fraudster leaves his/her voiceprint in the same way that our fingers leave fingerprints when we touch an object. This enables an organization to create and store voiceprints of known fraudsters. Each call that reaches the IVR or contact center can be proactively verified against this fraudster database. When fraudsters are identified, they can be denied access to prevent fraud from occurring.

9 http://news.softpedia.com/news/Australian-Bank-Call-Center-Staff-Easy-to-Socially-Engineer-177668.shtml

However, even in cases where fraudsters are successful, the voiceprint left at the crime scene can be used to identify and prosecute the criminal. This serves as a powerful deterrent to fraudsters, and leads to displacing fraudster activity to organizations that have less effective security measures such as PINs and agent handled security questions.

A voice is unique to the individual
A person's voice is unique, much like a person's fingerprint, iris or face. There are over one hundred voice characteristics that can be measured to determine who you are based on your voice. To identify a person using voice biometrics, a person's voice needs to be captured. This makes voice biometrics fundamentally different from knowledge based credentials, such as PINs and security questions. Someone cannot guess your voice, whereas someone can guess your PIN or answers to security questions. For the most part, a person's voice is not readily available on the internet, unlike the answers to security questions, such as a person's mother's maiden name. Although a recording of someone's voice can be captured by a malicious user, the voice is inherently static. Voice biometric systems can be dynamic, meaning they can be used to assess a person's identity during live conversations or by asking a caller to speak a random phrase.

Combined, these three key fundamental characteristics of voice biometrics lay the foundation for improved security within an IVR or contact center, and explain why Nuance customers have experienced up to a 10x reduction in fraud by migrating to voice biometrics.

10x

reported fraud reduction

What vulnerabilities does voice biometrics have?

As mentioned at the outset of this paper, no singly technology can eliminate all security vulnerabilities. Despite augmenting the security over PINs and security questions, fraudsters still have ways to compromise security. This section outlines the known vulnerabilities of voice biometrics.

Voiceprint re-enrollment (only applies to IVR)
This is a social engineering trick that can be leveraged by a malicious user, in the case where voice biometrics is deployed in the IVR. In this attack, the malicious user claims to the contact center agent that they are unable to authenticate with their voice, and that their voiceprint needs to be re-enrolled. If the agent complies, a fraudster can be enrolled in the system and be provided with access to a legitimate account.

This attack assumes that the fraudster is unknown to the organization, meaning that they have not called previously and been identified as a fraudster in the Nuance voice biometric system. Clearly, a known fraudster would not be able to conduct such an attack since the agent would be alerted.

Although similar to the PIN reset attack, the voiceprint reset attack is significantly easier to prevent. PINs get reset very frequently because they are forgotten. An agent will not question a user when they claim that they forgot a PIN. However, the need to re-enroll a caller in a voice biometric system is very rare. Voiceprint reset requests tend to occur at a rate of 0.2% of all enrollments. They also follow a common pattern. A caller enrolls in the voice

biometric system, and then is unable to successfully verify. This is often due to a problem that occurred during the enrollment process. As such, an agent can verify that the caller has recently enrolled and has not been able to verify. A caller that requests a voice biometric re-enrollment that has successfully authenticated previously is most likely either a fraudster, or does not need to be re-enrolled.

Due to the low frequency of re-enrollment requests, additional security procedures can be implemented compared to a PIN reset function, in order to minimize this attack by fraudsters.

In the case where a fraudster successfully compromises an account by enrolling their voiceprint, they have provided fraud investigators with biometric evidence at the scene of the crime. This evidence can then be used as evidence in court to assist prosecution of the fraudster. As such, the risk to the fraudster is immense, and in itself serves to deter such attacks.

It is important to note that this type of attack only applies to voice biometrics in the IVR. When voice biometrics is used in the contact center, in a live conversation with the agent, the caller is not aware that voice biometrics is being used, nor are they aware of the results of the voice biometrics assessment. As such, they cannot request to have their voiceprint re-enrolled.

Brute-force attack
A brute force attack on a voice biometric system, whether at the IVR or contact center, is a possible attack although it has a very low probability of success. In essence, this attack consists of a fraudster calling the IVR or call center numerous times until their voice is mistakenly accepted by the voice biometric system as belonging to a legitimate account holder. Vulnerability testing conducted on deployed voice biometric systems indicates that the rate of a success of a brute force attack is between 0.1% and 0.5%. This means, that on average, a fraudster would need to make between 200 and 1000 calls before they are able to access an account. The chances of a successful brute-force attack on a PIN based system are 20x to 100x higher than with voice biometrics.

There are several ways to mitigate even this lower level of risk. Nuance voice biometric solutions have the ability to detect cases where the same caller is dialing into multiple accounts. Such a caller can be blocked from accessing accounts. There may be legitimate reasons why a user calls into more than one account, for example if they mistakenly provide their account number incorrectly. As such, the three strikes rule can be applied. If the same caller dials into three separate accounts, then that caller will not be authenticated into any account.

The same rule can apply to a single account. If there are three concurrent failed authentication attempts on a single account, that account can be locked to minimize the probability of a successful attack.

Vhishing (only applies to IVR)

Vhishing is a vulnerability that afflicts both knowledge based factors, as well as voice biometrics. In the case of voice biometrics, the fraudster calls a legitimate account holder to persuade them to speak the passphrase used to authenticate to the IVR system, such as "At Sphere Bank, my voice is my password". The fraudster records the legitimate user speaking the passphrase, and can then play the recoding over the phone when prompted by the IVR to authenticate. To the voice biometric system, the caller is legitimate, as it is indeed the legitimate caller's voice that is being captured by the system. However, Nuance has several counter measures to eliminate or minimize the chances of a successful attack. These include playback detection and liveness detection. When used concurrently, they can reduce the chances of such a replay attack to between 0.5% and 2%.

Liveness detection entails performing a second biometric verification, using a random passphrase. This mitigates the above threat, as the fraudster will not have a recording of the legitimate account holder's voice speaking the random passphrase. Although a liveness check is impractical to conduct on all calls, it can be very beneficial prior to conducting high-risk transactions. More details on Liveness detection can be found in the Nuance white paper titled "Mitigating Recording Threats with Nuance VocalPassword".

Liveness detection is most effective when paired with playback detection. Playback detection is a patented Nuance technology that compares the utterance of a passphrase with past verifications for a particular account. This enables the detection of a recording.

As with the voiceprint re-enrollment attack, Vhishing can only be used to attempt an attack on a voice biometric system securing access to an IVR. In the contact center application, voice biometrics monitors a live conversation with the agent, and as such an attack using voice recordings is not feasible.

Summary of vulnerabilities by authentication method

The table below provides a comparative summary of the various security vulnerabilities for PIN, security questions and voice biometrics. Each is rated based on the success rate that a malicious user can expect to achieve by exploiting the vulnerability. A rating of high was assigned to a vulnerability that could successfully be exploited more than 25% of the time. A rating of medium was assigned to a vulnerability that could be successfully exploited more than 5% of the time, but less than 25%. Finally, a rating of low was assigned to vulnerabilities that could be exploited at a rate of less than 5%.

Security Vulnerability	PIN	Security Questions	Voice Biometrics
Brute Force Attack	Medium 10%+ success rate	N/A	Low 0.1% to 0.5% success rate
Credential Sharing	High 100% success rate	N/A	Low 0.5% to 2% success rate
Hacking	Low	Low	None 0% success rate
Phishing	High 72% success rate	High 72% success rate	N/A
Vhishing	Medium	Medium	Low 0.5% to 2% success rate
Credential Reset	High	N/A	Low
Internet Search	N/A	High	N/A
Social Engineering	N/A	High 67% success rate	N/A

This table clearly illustrates that when voice biometrics is compared to PINs and agent handled security questions, voice biometrics delivers a considerable improvement to security. This highlights why financial institutions that have deployed voice biometrics in their IVR or contact center have experienced reductions in fraud that have reached up to 10x.

About Nuance Voice Biometrics
Nuance is the global leader in voice biometric solutions, with over 35 million enrolled voiceprints and a global customer base that spans all major industries. Nuance has developed over the last 12 years unrivaled experience in delivering successful voice biometric solutions that enable organizations to improve customer satisfaction reduce costs and improve security.

About Nuance Communications, Inc.
Nuance Communications is reinventing the relationship between people and technology. Through its voice and language offerings, the company is creating a more human conversation with the many systems, devices, electronics, apps and services around us. Every day, millions of people and thousands of businesses experience Nuance through intelligent systems that can listen, understand, learn and adapt to your life and your work. For more information, please visit nuance.com.