

# Fraud prevention best practice toolkit.

Practical steps to create and maintain your optimum fraud prevention solution.

# Table of contents

- 1 Characterising your organisation's attitude to fraud / p2**
- 2 Assessing your fraud landscape / p4**
- 3 Weighing the benefits of offline and real-time fraud detection / p5**
  - Offline fraud detection
  - Real-time fraud detection
- 4 Understanding the fraud attack cycle / p7**
- 5 Maximising the value of alerts and management information / p8**
  - Optimising blacklist management
  - Analysing your data and prioritising alerts
  - Capturing the richest possible metadata
- 6 Evolving your defences in line with changing fraud behaviour / p10**

## Why traditional identity proofing is dead

Personal identity numbers (PINs), passwords and complex security questions - the answers to which many customers find impossible to remember - have finally had their day. A constant barrage of data breaches has led to widespread compromise of personal data and user credentials - so much so that according to one Gartner analyst<sup>1</sup>, an average of 15-30% of customers fail identity proofing, while up to 60% of criminals pass them.

It's hard to imagine any other aspect of corporate activity where a 15-30% failure rate would be acceptable - let alone one with such a negative impact on the customer experience, or that criminals find so easy to exploit. To address this, it's estimated that within five years, there will be no more knowledge-based authentication to secure sensitive accounts or information that can be targeted for fraud.

One reason is the increasing number of enterprises that are turning to solutions such as Nuance Security Suite to replace PINs, passwords and security questions with voice biometrics and other characteristics that are much harder for fraudsters to compromise. In 2017, 90 billion knowledge-based credentials were maintained worldwide, compared to the 300 million voiceprints used for authentication with Nuance voice biometrics solutions. But all the indications through the use of these solutions in the marketplace are that the number of voiceprints in active use is growing exponentially, with the total doubling in 2017 and a similar trend in 2018.

So if yours is one of the many leading businesses that are rethinking their approach to customer authentication and fraud prevention within their customer care channels, this best practice toolkit will enable you to benefit from Nuance's extensive experience in helping organisations create and maintain their optimum fraud prevention solution, including due consideration of key aspects such as:

- Characterising your organisation's attitude to fraud
- Assessing your fraud landscape
- Weighing the benefits of offline and real-time fraud detection
- Understanding the fraud attack cycle
- Maximising the value of alerts and management information - including optimising blacklist management, analysing your data and prioritising alerts, and capturing the richest possible metadata
- Evolving your defences in line with changing fraud behaviour

## Characterising your organisation's attitude to fraud

It might seem obvious that the goal of any fraud prevention solution would be to minimise an enterprise's losses due to fraud. But in reality, different organisations may have widely differing attitudes to the levels of fraud they are willing to accept.

A newly established credit card company, for example, might have the aim of driving transaction volumes and revenues as a means of gaining a foothold in the market.

Top 3 benefits delivered by voice biometrics:



Reduced costs



Decreased fraud



Brand differentiation

<sup>1</sup> Avivah Litan, Gartner Analyst, Absolute Identity Proofing is Dead, November 2015

It may therefore want to make it as easy as possible for customers to spend their money using its credit cards, and not want to do anything that impedes the customers' ability to spend.

This might mean that rather than requiring even a basic security mechanism such as a PIN, it allows transactions to be verified by the customer's signature – even though this is rarely checked and offers virtually no defence against fraud. In effect, the company will have decided that it is prepared to write off significant fraud losses to enable customers to spend as much as possible, because that spending is where it sees its revenues being generated.

Contrast this with the attitude of a financial institution that is also offering a new credit card, but may have come from a background of providing pensions, mortgages, insurance or other long-term financial products; or whose agility is limited by legacy systems and processes, or a more conservative culture and mindset. Because this company is used to dealing with one-off, high value transactions, rather than the high volume, lower value transactions associated with credit card purchases, it may have a much lower tolerance for risk – including fraud.

As a result, its fraud prevention mechanisms, processes and controls around authentication and verification - even for something as simple as issuing a new PIN - may be much tighter and difficult to navigate than a more innovative credit card company's. Its corporate culture may also make it very difficult to change this approach, despite the fact that this will have a negative and potentially damaging impact on the customer experience.

So even though both these companies could be offering otherwise identical credit card products, their dramatically differing perspectives in terms of their attitude to risk will have significant effects not just on their tolerance for fraud, but on everything from their authentication and verification processes to their communications and dealings with customers.



### **Nuance best practice**

These examples illustrate two extremes of the balance every organisation needs to strike between security and the customer experience. On the one hand, if you implement extremely tight security controls, this will typically be at the expense of creating a cumbersome and unfriendly customer experience that can negatively affect your brand. On the other, if you want to make the customer experience as simple as possible, this may leave gaps in your defences for which you'll have to be prepared to accept a certain level of losses due to fraud.

To help you achieve the best solution for your business, when Nuance undertakes a risk assessment prior to solution deployment, one of the aims is to establish your organisation's appetite for risk, and where on the spectrum between convenience and security you are prepared to set the bar. A best practice is therefore to aim for the point between these two extremes that achieves the optimum balance between minimising fraud losses while not impinging your customer service. Nuance then enables you to achieve this through:

- Voice biometrics authentication and verification to ensure your customers are who they say they are, and so improve the security of your 'front door'.
- Counter-fraud measures that enable you to identify known and persistent fraudsters and strengthen your defences against them.

## Assessing your fraud landscape

Before you start to reduce your vulnerability to fraud, you also need to establish a clear picture of your 'fraud landscape', including the levels of fraud and the nature of the attacks you're currently experiencing.

This will enable you to sanity check the scale of your fraud problem; whether the cost of fraud is acceptable given the nature of your business; how and where Nuance can best help you to reduce the fraud attacks and losses you're suffering; and the implications of your risk appetite when it comes to striking the right balance between security and quality of the customer experience.

As an example of how this works in practice, the fraud landscapes of a utility company, a bank and a credit card company will be completely different, as will the counter-fraud measures they will need to put in place.

- When a utility such as a cable or satellite TV company acquires a new customer, their liability will generally be limited to the value of the set-top box (and satellite dish) installed at the customer's premises, the cost of which will be low and will be recouped over the customer's contract. Potential fraud losses are likely to be related to process abuse, such as using a domestic set-top box on commercial premises, or attempting to break the box's encryption to access channels that have not been subscribed to.
- For a bank, however, not only will the financial liability be much larger, but the nature of the fraud attacks being experienced will be different, as these will be mainly transactional. If a customer suffers an account takeover attack, for example, the fraudster might take out a loan against the customer's name, use their overdraft facility etc. – for all of which the bank will be liable to reimburse the customer.
- Similarly, if a credit card company offers a fraudulent customer a sizeable credit limit which the fraudster uses and never repays, all of this money will have been lost by the company.

Based on these loss profiles, a utility might therefore put in place significantly fewer fraud checks and hurdles to purchase than a bank or a credit card company would require, as this makes sound business sense given the differences in potential losses.

### Nuance best practice

Many organisations have a wildly inaccurate understanding of their fraud landscape, and so find it very difficult to make fraud-related decisions based on sound financial information. While some financial institutions, for example, can experience annual fraud losses amounting to many millions of pounds or euros, others have estimated these losses to be in the tens of thousands – the implication being that either they were not experiencing any significant fraud, or, more likely, they were blind to the fraud that was occurring.

Even those enterprises that have a relatively good grasp of their fraud landscape can find that, having implemented Nuance's FraudMiner™ counter-fraud solution, they had significantly underestimated the full scale of the losses they were experiencing, as FraudMiner often uncovers a variety of attacks that were previously undetected.

## Weighing the benefits of offline and real-time fraud detection

Having established your appetite for risk, and the nature of the fraud landscape to which your business is exposed, you'll be in a position to make better informed decisions about the benefits you can expect from offline and/or real-time fraud detection.

### Offline fraud detection

Many enterprises report they have achieved a very rapid reduction in fraud losses following the introduction of offline fraud detection – with one organisation preventing £4.5 million of losses in its first six weeks using Nuance Security Suite in its contact centre.

A credit card company, for example, may be experiencing problems with account takeover attacks, such as fraudsters requesting replacement cards to be sent to a different address to the genuine account holders'. Issuing a new card involves manufacturing, printing and delivering the card – a process that can typically take several days. But because of this natural delay, there is time within the process for the attack to be detected and the new card cancelled before it is dispatched. Offline fraud detection is therefore highly effective in stopping these types of attacks.

### Real-time fraud detection

By comparison, if a bank enables customers to transfer money from one account to another using a faster payment method, this can be abused by a fraudster making a call to the bank's contact centre, and without real-time fraud detection the money involved may be lost before the end of the call.

Anecdotal evidence provided by one Nuance customer suggests that in these situations, the money transferred can be recovered in around 50% of cases (for example if it has been transferred to another bank within the same group). Real-time fraud detection enables losses from the other 50% of frauds that cannot be reclaimed (for example because the money was transferred to a different country) to be stopped.

- A key difference between offline and real-time fraud detection is that the latter is truly customer-impacting, because the caller is on the phone when the fraud assessment is being made. So unlike offline detection where a call can be investigated after it has been completed, if a genuine customer is flagged as a fraudster, there is every chance they may become frustrated or upset.
- With offline fraud detection, the entire duration of a call is checked, whereas with real-time detection this can be as early as during the identification and verification process. It can, however, be beneficial to delay fraud alerts until as late in the call as possible - for example, just as the agent is about to commit a payment - as this will enable the maximum amount of audio to be captured on which to base the fraud assessment.

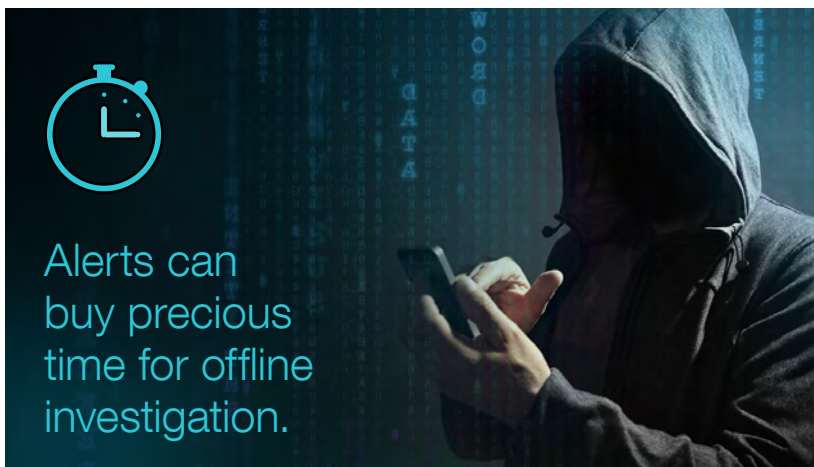
The way contact centre agents handle real-time alerts also needs careful planning. The nature of real-time alerting is that statistically there will always be more false alerts than true alerts, because there are far more genuine customers than fraudsters phoning in. This means that if an alert is raised it should not automatically be assumed to be a fraudulent call.

Agents need to be clear about how to react to such alerts – for example by asking more security questions or transferring the call to a different team – and this process may vary according to the organisation's risk profile, but also the transaction that is the reason for the call.



One organisation **prevented**  
**£4.5 million** of losses in its first six  
weeks using Nuance Security Suite.

- If, for example, someone is calling to ask for a replacement bank card and this has flagged an alert, while it is a potentially risky transaction, it is also one that can be dealt with offline due to the time delay in issuing the new card. It may be that the process for such an alert is therefore to tell the caller the new card will be delivered within the next five working days, but then investigate the call offline to determine if it really was a fraud attack.
- By comparison, if somebody is trying to move a large amount of money to a new payee that has just been set up and this flags an alert, a very different reaction would be required. The process could be that rather than making the payment immediately, the agent tells the caller that payment will take place within two hours subject to the usual fraud checks, and this is then put into a pending queue for offline investigation during the two-hour window.



At Eastern Bank,

**94%**

of agents reported voice biometrics makes it easier to deliver quality service, and

**60%**

that their job satisfaction has improved since deploying voice biometrics

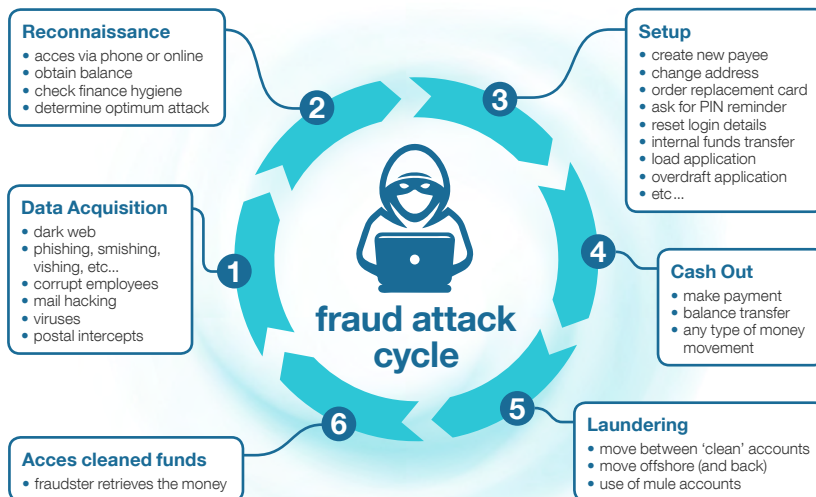
### Nuance best practice

So which solution, or combination of solutions, is going to be right for you?

- Understanding the nature of your fraud landscape will allow you to determine whether improving your ‘front door’ security through voice biometrics authentication will help, and if so how much. This will also tell you whether you require real-time fraud detection, or if the fraud you’re experiencing can effectively be managed offline.
- If analysis of your fraud landscape determines, for example, that you could stop 40% of fraud with an offline solution, this is typically the easiest to deploy, requires the least integration and development effort, and could be up and running within just three months. You might therefore decide to implement an offline solution to give an immediate reduction in fraud, while also starting to look at rolling out real-time detection and authentication to tackle the remaining 60% of fraud, and so establish a strategic deployment plan.
- For real-time fraud detection, the management of and reaction to fraud alerts will need to be different to that for offline, and this may depend on both the nature of the transaction and your risk profile.

## Understanding the fraud attack cycle

One of the reasons why fraud benefits can be realised very quickly with Nuance Security Suite is that fraudsters tend not to perform a single attack on an organisation, but do so repeatedly, with each of their attacks having multiple steps. Another way of thinking about how best to reduce your fraud losses is therefore through this 'fraud attack cycle' of data acquisition, reconnaissance, setup, cash out, laundering and access to cleaned funds.



**FraudMiner** can detect fraud at steps 2, 3, and 4.

Detecting fraud at steps 2 or 3 prevents it from reaching stage 4. Stage 4 is where real-time detection becomes important. Offline can be used at stages 2 and 3, minimising the need for real-time at stage 4.

**That is, we can use offline to stop the fraud ever reaching the cashing out stage.**

Steps 5 and 6 are likely to be with different banks to the victim. If the banks are using **AudioShare** the true identity of the fraudster can be found with **FraudMiner**.

- Most organisations are aware that in terms of data acquisition, some of their data may have been hacked and made available on the dark web; some of their customers will be socially engineered into giving up their details; they may have rogue employees selling customer data; etc.
- When a fraudster has acquired this data, there will usually be a reconnaissance step where they will use it to start accessing an individual's account, uncover details such as the account balance, and potentially look at some kind of account history to help them perpetrate an attack in an intelligent way.
- But because most organisations' focus is typically on events in the cash out phase (i.e. the point at which they lose money) which can be easily measured and tracked, they may not have considered the impact of reconnaissance and setup activities that do not immediately result in cashing out.

### Nuance best practice

FraudMiner can detect fraud during the reconnaissance, setup and cash out phases; and offline detection can also be used during reconnaissance and setup, thereby reducing the need for real-time detection at the point of cashing out. So while many banks, for example, have tended to focus heavily on cashing out, through FraudMiner they are finding they're receiving significantly more reconnaissance and setup calls than they had ever anticipated, and so are able to act on these much earlier in the fraud attack cycle.

- According to the experience of one major bank, around 25% of fraud alerts generated by FraudMiner are in the pre-cashing out phase, meaning they are able to identify that an account has become compromised, take preventative action very early (and before there is any threat of money being lost), and so stop the fraudsters in their tracks.

More and more fraudsters are successfully being prosecuted where voice biometrics is used as evidence



- This can have the additional benefit of making the bank appear extremely conscientious to its customers, by warning them their account and security has been compromised, and recommending they check the status of any accounts with other organisations secured using the same credentials.

Even after the cash out phase, enterprises signed up to Nuance's AudioShare service can use FraudMiner to establish the true identity of known fraudsters that are in the process of laundering or trying to access cleaned funds – which in turn opens up the possibility for those fraudsters to be located, arrested and prosecuted through law enforcement.



## Maximising the value of alerts and management information

By identifying a single fraudster's voice on the phone channel, one bank was able to identify 1,390 accounts that had been compromised online. This highlights the vital importance of continually analysing the alerts and management information (MI) being generated by FraudMiner.

While in the early stages the focus of data gathering is to enhance understanding of the fraud landscape, once offline and/or real-time fraud detection has been implemented, this then switches to MI dealing with the specific threats to which the business is exposed.

### Optimising blacklist management

To make their MI actionable, most businesses maintain a variety of lists, including:

- Blacklists (also known as fraudster watch lists) - which always refer to known fraudsters. Blacklist entries are added based on them being defined as 'confirmed fraud', that is the bank can track back that individual to a call where the bank have confirmed a fraud loss
- Watch lists - a term used interchangeably with blacklists, which can also be used for non-fraud reasons such as highlighting vulnerable persons, customers who have had issues with their accounts, or customers who are repeatedly accused of fraud as a result of generating blacklist matches.
- Whitelists - which refer to customers who, even though they are triggering fraud alerts, are known to be trustworthy.

Because a very high percentage of fraud losses are perpetrated or caused by a very small number of fraudsters, blacklist management is the 'crown jewels' of fraud detection, and where its value lies.

- Once an organisation has compiled a list of known fraudsters, it will want to do as good a job of matching against this list as possible, by accumulating more and more data about each fraudster.
- Each time the fraudster calls, this will improve the performance of their voiceprint, meaning it will score better when the fraudster tries to attack. Through this, the enterprise will not only get a better match and more confident alert; it will also reduce the likelihood of false alerts with people with similar voices to the fraudster.

---

By identifying a single fraudster's voice on the phone channel, one bank was able to **identify 1,390 accounts** that had been compromised online.

---

The key to blacklist management is improving the separation of the fraudster's voice from non-fraudsters' voices; improving matching against each fraudster's voice; and improving performance of each individual voice on the blacklist.

- Whether you have 100 or 1,000 entries in your blacklist, this will generate a certain number of false alerts, and if your blacklist is 10x that size, you will probably generate 10x as many false alerts. So if possible, you should try to keep the size of your blacklist down to avoid generating more and more false alerts.
- If you are genuinely under attack by 1,000 fraudsters, then all of those people should be on your blacklist. But if your MI shows that 200 of those have not attacked you in the last year, these should maybe be retired (but not deleted) from the list, as all they are doing is raising false alerts.
- You should though also bear in mind that retired fraudsters may come out of retirement, and periodically run retired entries against your data to see if they have come back into operation, and if so put them back on your blacklist.

Analysis of blacklist performance; determining which entries to keep and which to retire; and making sure your blacklist entries are as relevant and as good performing as possible; are all vital to the operation of your fraud detection solution. But from a customer experience point of view, so is how you handle customers who get added to your blacklist despite having no connection to fraud.

- If, for example, you have a small group of customers with a particularly distinctive ethnic or regional accent, and your counter-fraud solution identifies a fraudster (or fraud gang) with the same accent, the nature of the biometric engine and background model means other totally innocent members of the group may themselves start triggering regular fraud alerts.
- You can address this by adding the genuine customers' voiceprints to a whitelist, so that if their voice matches the blacklist but also the whitelist, it does not raise an alert.

#### **Analysing your data and prioritising alerts**

As well as blacklisting management, the team tasked with managing your counter-fraud solution will need to be responsible for data analysis and alerts. You therefore need to structure your team accordingly, with one group managing your blacklist, another working with and learning from your MI, and a third group prioritising alerts.

MI analysis includes analysing alerts, fraudster and voiceprint performance, alerting thresholds, and extrapolating to other channels via metadata. Prioritising alerts involves working the most 'fraud-rich' alerts first, and prioritising by score, calling line identity (CLI), line of business, etc.

- Organisations implementing FraudMiner report that the data they are receiving from the solution is more fraud-rich than their other alerting mechanisms, and they are identifying more fraudsters as a result. They may therefore reallocate people from other alerting systems onto FraudMiner to handle alert investigations, as these are the most productive.
- Enterprises may also find that because of FraudMiner, the number of alerts and fraudsters being detected may decline over time, as the business is no longer an easy target, fraudsters are being less successful, and so they focus their efforts elsewhere. The number of people required to investigate these alerts may therefore flex over time.
- How effectively the team operates will be determined by the MI assessments, looking at the fraud attack rates, and analysing the behaviour of the fraudsters being revealed through MI. This means that ideally the whole team will work in harmony – managing the blacklist and also working the alerts based on what their MI is telling them.

### **Capturing the richest possible metadata**

All of the team's activities will be strengthened by capturing as much metadata as possible when the counter-fraud solution is being developed and built, as this enables a more detailed analysis of the fraud landscape.

Rich metadata can also be used to filter calls for backward searches and clustering, gain insight into fraudster behaviour, identify process weaknesses and simplify retrieval of original calls. As just some examples of the benefits of this:

- Many organisations operate numerous different phone numbers offering multiple routes through which customers can contact them. So, for example, a bank might have different phone numbers for customers calling about mortgages, credit cards, business banking etc. By capturing this information it may be possible to identify certain routes into the business that are more susceptible to fraud attacks, and so might have a process vulnerability that fraudsters are exploiting.
- Through this, one bank found that fraudsters were attacking its Fraud team, because the number of checks the team were making was less than the main Banking team. A fraudster would call the Fraud number and say they had received a voicemail that there had been a fraudulent payment on their credit card, but that this was a genuine payment. In reality there had been no outbound call to the fraudster so there was no record of it, but the Fraud team assumed the agent had failed to make a note of the call, and because the fraudster had passed the security check they lifted the fraud marker on the account. The Fraud team were therefore being socially engineered by the fraudster to open up the account.
- Another bank uncovered a process flaw with its Collections team, whose primary concern was to collect late payments on credit card accounts, rather than where the money to do so originated. To exploit this, fraudsters would take out a card with the bank, spend up to their credit limit and then default on the account, prompting a call from the Collections team. Although customers cannot pay off one credit card with another, the Collections team were not following this rule. Fraudsters were therefore able to service their debt using another fraudulently obtained card, extending the life of both cards.

## **Evolving your defences in line with changing fraud behaviour**

The Nuance Security Suite delivers omnichannel security and fraud prevention across digital, telephony and self-service channels; with voice biometrics authentication and counter-fraud solutions that reduce organisations' vulnerability to fraud, and stop fraudsters in their tracks.

However, there is no single capability that will effectively detect fraudsters in all situations, which is why fraud prevention, like authentication, requires a layered approach including the application of AI to detect fraud patterns and improve the detection of fraudsters as they try to attack customer care channels.

Also, as consumers and enterprises increasingly adopt biometrics for authentication purposes, fraudsters will inevitably divert their attention to trying to compromise the biometrics involved, which is why Nuance is developing a range of technologies to mitigate and prevent these attacks.

- For fraudsters using spoofing mechanisms such as recording attacks, synthetic voice attacks for voice biometrics, using a picture or video to try to compromise facial recognition, etc. Nuance is developing sophisticated anti-spoofing algorithms in areas such as liveness, channel playback, footprint playback, synthetic speech, picture and video detection, and ANI spoofing detection.

- To enhance the performance of voice biometrics and text-based communications, Nuance ConversationPrint™ adds behavioural biometrics based on a caller's vocabulary, grammar and sentence structure, including scripts. This can be used to identify individuals or fraud gangs through their speech patterns and habits, provide another authentication layer to complement other biometrics, detect net-new fraudsters and reduce false fraud alerts.
- To enable further analysis of phone calls beyond the caller's voice, Nuance Channel ID can detect a caller's phone type, while GeoID can identify the country and city the device is associated with via the phone network. This, along with Nuance Device Print (which verifies the current device used against devices used by the legitimate account holder in the past), can then be used to add to the metadata associated with particular individuals, complementing authentication and assisting with the assessment of risk.
- Behavioural biometrics can also track, measure and analyse a customer's or individual's interaction with their device, as an additional means of enabling them to be identified.
- All this data can be fed into an AI engine to holistically review the results; understand the weights and values of the various data points, behavioural biometrics and contextual factors; and produce a unified authentication score and unified fraud score, per interaction.

To summarise, just as behavioural prints can be created for legitimate customers, and used as a credential or replacement for PINs or passwords; so they can also be generated for fraudsters, to detect and identify known fraud behaviours and profiles, and prevent them compromising customer care channels.

To find out more about what Nuance fraud prevention solutions could be doing for your enterprise, please contact [Brent Hunt](#).



---

#### 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](http://nuance.com).

---