White Paper

Smarter Communications Surveillance with AI

New technology enabling financial services firms to improve their tools and reduce risk

EXECUTIVE SUMMARY

Communications in business are rapidly changing, due to the proliferation of channels and an evolution of the way people work. To protect financial services firms from increased risk, communications surveillance tools need to change, too.

Traditional lexicon-based methods of communications surveillance are no longer adequate to monitor employee misconduct. Though still useful, these tools are only powerful enough to detect the use of predetermined language. They lack the ability to understand context, or the severity of threats. As a result, traditional lexicon-based models produce an overabundance of alerts that bog down compliance teams and allow real threats to get lost in the clutter.

Artificial Intelligence (AI) has emerged as a solution that can provide essential context to alerts generated by lexicon-based models. This is known as **Natural Language Processing (NLP)**. Employing **AI** in the form of **NLP**, in conjunction with traditional methods, improves detection accuracy and enables the identification of genuine misconduct.

In this White Paper, we discuss limitations of traditional communications surveillance and explore how **AI** and **NLP** empower financial services firms to enhance their communications surveillance tools and mitigate risk effectively. Additionally, we provide insights on implementing **NLP**-enabled lexicon models while ensuring compliance and preserving a seamless user experience.

TABLE OF CONTENTS

1. INTRODUCTION

In financial services, communications are how business gets done. But communications also create risk. Misconduct can hide in the thousands of calls, texts, and emails that regulated employees engage in every day.

Misconduct – intentional or not – can erode customer trust, result in regulatory fines and sanctions, and negatively impact a firm's reputation.

Without effective surveillance tools, communications are a liability. The context needed to understand what regulated employees are saying and doing will stay hidden and concealed.

The fact is, smarter approaches to communications surveillance are becoming more essential by the day, due to the following factors and trends.

Today's Challenges

Remote/hybrid work, and changing work practices

According to **Deloitte's** *Market Abuse Outlook 2022*[1], the COVID-19 pandemic fundamentally changed the way we work and live. Hybrid work is becoming the norm and this is increasing the risk of market abuse behaviors. There is a greater reliance on communications surveillance when employees aren't physically in an office where it's easier to monitor them.

> *"Every organization should prioritize this issue. Employee misconduct poses one of the most expensive risks for organizations worldwide. The rise in employee misconduct is the result of a confluence of factors."*
>
> *Jonathan Frieder,*
> Principal Director, Accenture

Evolving communication channels

The risk of misconduct outside of management's purview is also increasing. Growing use of unapproved communications devices, such as personal mobile phones, allow regulated employees to circumvent controls.

Adding to this is the increased adoption of digital channels (like SMS-based texting and email, *WhatsApp*, *WeChat*, and even unified communication platforms like *Zoom* and *Microsoft Teams*). New channel adoption is contributing to information overload, and increasing the volume and variety of communications that need to be surveilled.

Still, it's estimated that 60% of firms are not yet monitoring newer channels such as *Microsoft Teams*, *Bloomberg*, *WhatsApp*, *Slack*, *Telegram* and *Signal*.

Incidents are increasing

As mentioned above, employee misconduct remains one of the costliest sources of risk globally, driving the need for increased communications and holistic surveillance. According to the *2022 Ponemon Institute Cost of Insider Threats Global Report*, incidents of insider crime increased 44% between 2020 and 2022, with the average cost per incident increasing by 33% to $15.38 million.

Fifty-six percent of these incidents were due to negligence (e.g. employees falling prey to phishing and oversharing of data). The same study revealed that 74% of surveyed firms reported incidents where insiders emailed sensitive data (without proper authorization) to outside parties.

Approximately 67% of firms included in the 2022 study reported experiencing between 21 and 40 incidents per year, with 2023 year-to-date results trending even higher.

1 https://www2.deloitte.com/content/dam/Deloitte/us/Documents/financial-services/us-fsi-market-abuse-outlook-oct-2022.pdf

Excessive false alerts

Greater volumes of alerts are in turn generating higher volumes of false positives, and increasing surveillance costs.

According to a **PWC** *Market Abuse Surveillance Survey*, over a period of twelve months, seventeen of the banks participating in the survey raised a combined global total of 40 million trade and eComms alerts. The study revealed a false positive population of 99.99%, meaning that only a fraction of alerts were indicators of true risk.

A separate **Chartis** research survey, *The Future of Trader Surveillance*, confirms that one of the biggest pain points plaguing financial compliance and risk teams is the high number of false alerts. In fact, of the average 1,000 to 1,500 alerts that tier one banks generate daily, approximately 99% turn out to be false.

Even without factoring in false alerts, communications surveillance can be a very time-intensive and manual process. Still, daily false alerts and associated labor costs involved in dispositioning and investigating them can be a huge expense, costing firms millions of dollars annually.

Cost of false alerts - average tier 2 bank

		Cost per alert or case	Total daily cost	Total annual costs
# Daily false alerts without a case	800	$10	$8,000	$2,080,000
# Daily false alerts resulting in a case	4	$500	$2,000	$520,000
			$10,000	$2,600,000

"The regulatory focus on off-channel communications, particularly business-related interactions conducted via unauthorized communication channels like personal devices, is intensifying. This presents a significant challenge for organizations aiming to enforce communication policies and meet regulatory requirements related to records retention."

Jonathan Frieder,
Principal Director, Accenture

"There is a clear need to integrate communication surveillance into your risk management program. Monitoring and analyzing employee communications, encompassing emails, instant messages, and phone calls, are vital measures to identify potential misconduct or unethical behavior that violates professional standards, code of conduct or ethical guidelines."

Jonathan Frieder,
Principal Director, Accenture

More Fines

According to **Deloitte**'s *Market Abuse Outlook 2022* with the increased use of new communication platforms, regulators are taking notice, and especially scrutinizing the use of unapproved channels. The report goes on to say, "In the Americas, the SEC and CFTC have imposed notable fines of USD 200 M (per firm) related to business communications on personal devices via communication applications that were not being appropriately captured or surveilled. And regulators are expanding these types of enforcement actions and fines to European banks as well."

In a separate development, a leading international bank was hit with a two-hundred million dollar fine for record-keeping and surveillance lapses related to regulated employee use of an unapproved communications channel (WhatsApp).

In this White Paper we will examine how financial services firms are managing communications surveillance today, and why the above market trends demand a new approach that supplements traditional lexicon-based surveillance with **Artificial Intelligence** (**AI**) and in particular **Natural Language Processing** (**NLP**).

We'll also share tips and best practices on how to apply these complementary technologies to reduce false alerts and compliance costs, and identify misconduct across all of your communication channels.

2. HOW COMMUNICATIONS ARE MONITORED TODAY

When asked how firms are conducting communications surveillance today, *Jonathan Frieder*, Principal Director for **Accenture**'s U.S. Regulatory and Compliance Practice, says the answer falls in one of three categories. Based on his first-hand experience meeting and working with clients in the financial services industry, Frieder says, firms either:

1. Have nothing in place currently (e.g. they may not be capturing certain types of communications at all, or if they are, they're not surveilling them)
2. Are using older, less-advanced technology (typically lexicon-based solutions models and algorithms that can create challenges for contextual understanding, and require regular manual tuning)
3. Are utilizing a third-party solution or looking to create their own solution in-house

This view is also supported by Frieder's regular discussions with third party technology solution providers about trends they are collectively seeing in the industry.

According to JWG's ***The State of Holistic Trade Surveillance Benchmark Research Report***[2], the majority of firms [55%] performing communications surveillance are using manual processes to enhance their use of lexicon-based search criteria.

An equal number [18%] are relying on key words and taxonomies, which can assist with more advanced detection. However, no firms indicated that they are taking an ontological approach that would help contextualize issues and identify outliers.

This low-tech approach to semantic technology appears to be directly linked to the way alerts are managed. Seventy-five percent of firms indicated that they use rules to eliminate false positives. Interestingly, only 34% indicated that compliance is enabled to define models on the fly to examine behaviors and receive alerts.

2 Download a copy of the report here: https://actimize.nice.com/JWG-Benchmark-Research-Report

3. THE CHALLENGES OF LEXICON-BASED SOLUTIONS

According to *vocabulary.com*, the term lexicon refers to "the total stock of words and word elements that carry meaning." Essentially, it's a taxonomy or list of text or words that would be considered meaningful.

Compliance-driven organizations use lexicons during monitoring of electronic and audio communications to try to identify certain types of actions or behaviors that might signal misconduct and market abuse. Lexicons, in this instance, are essentially a taxonomy or list of keywords or phrases that might hold meaning. For example, the term "risk-free" might indicate a registered representative making an inappropriate promise, or the sentence "let's keep this a secret" might reveal insider dealing.

Lexicon-based solutions focus on finding words rather than understanding the true context of communications.

Financial services firms use lexicons in two ways:

1. to create lexicon-triggered alerts to identify communications to be reviewed by a compliance analyst, and
2. to perform manual searches on communications during active investigations.

In order to use either of these methodologies with voice communications, the communications need to be accurately transcribed before being fed into any monitoring tool.

Lexicons are widely utilized because they are easy to explain to regulators, simple to use, and easy to manage. But lexicons also come with their own set of challenges.

The challenges of using lexicons alone

Traditional lexicon-based communication surveillance techniques offer flexibility and ease of use, but come with many challenges:

Lexicons focus on finding words rather than understanding the true context of communications

Lexicon-based surveillance relies on the presence of specific keywords or phrases to identify risky communications. However, by nature, language is highly contextual and ambiguous, and certain words or phrases can have different meanings depending on the context of the communication. As a result, innocent discussions can be misinterpreted as harmful or suspicious.

Lexicons identify key words, but absent intent or sentiment

Lexicons can identify when something is said, but not the intent behind the words. For this reason, lexicon-based solutions aren't always one-hundred percent effective in identifying misconduct.

For example, a lexicon-based surveillance solution might flag communications containing the words 'Don't tell anyone...' based on inferred intent, but not all would be flagged correctly.

"Don't tell anyone that I am buying Microsoft shares." (correctly flagged)

"Don't tell anyone that I am buying toys for my kid." (incorrectly flagged)

Also, while lexicon-based approaches have their advantages in simplicity and ease of implementation, they may not be sufficient for more complex tasks that require understanding context, emotions (such as sarcasm), or implicit meanings.

Lexicons find too much

In line with above, lexicon-based surveillance systems, which are created to alert on every communication containing certain key words, will do so regardless of the context in which the words are used. Due to the high number of alerts triggered, firms are often tempted to either pare down the input parameters or resort to only reviewing a sample of flagged communications. This haphazard approach leaves firms open to risk.

Lexicon-based alerting can create excessive noise

Alerting on specific lexicons may seem simple and straightforward but that very simple approach can create problems, too, especially when it comes to eComms surveillance.

Most emails today contain standard disclosure text. When emails are ingested into the communications surveillance solution, alerts are often generated based on this irrelevant content.

And the problem is magnified as emails are forwarded, because every time the email is regenerated, it creates more opportunities for more false alerts. For example, if a single email was forwarded ten times it could result in ten false alerts.

The only way to alleviate this problem is to manually add new disclosure language into the lexicon dictionary (word for word) as an exception, so that every time new disclosure text is received, it doesn't trigger future false alerts. This can be very time consuming for compliance teams, as disclaimer content can vary by company, geography, and language.

Lower accuracy and quality of alerts

Lexicon-based surveillance is also known to generate higher numbers of lower quality alerts (false positives and negatives), which can cause compliance analysts to spend inordinate amounts of time on non-productive activities.

Difficulty 'understanding' industry slang/ trader jargon

As a predefined list of words, lexicons may not cover all the possible word variations, slang, or jargon used by traders in financial domains. Lexicons can easily become outdated and require frequent manual curation as new terms and phrases evolve.

Additionally, as regulated employees embrace new communication modalities and devices, lexicon-based surveillance can have a difficult time keeping up with language nuances and constantly evolving lingo. Consider for a moment text abbreviations like TLDR, LMAO, SSDD and FCIG that were unheard of just a few years ago. Languages and dialects can also complicate lexicon surveillance.

Simple evasion detection

Individuals with malicious intentions can always find ways to bypass lexicon-based surveillance systems by using coded language, misspellings, abbreviations, or synonyms that are not included in the keyword lists.

For example, a lexicon rule could be designed to look for the word "collusion," but it could be bypassed by using terms like "Let's double team it."

Because lexicons don't look for the overall context or understand the complete interaction, they are easier to circumvent. When you create a lexicon dictionary, you are essentially publishing the rules. Like any other system, if you know the rules, they're easier to evade.

Inability to adapt to changing behaviors and new risk profiles

The only certainty in communications compliance is that things will never stay the same. Regulated employees will adopt new communication modalities and devices, and may even change what they say and how they speak to circumvent controls.

For example, employees may switch to WhatsApp from SMS or other messaging platforms, change the way they talk, or switch to personal devices from firm issued devices.

In this constantly changing environment, detecting misconduct can be a little like trying to hit a moving target. First, the new platforms might not even be monitored, and even if they are, lexicons are based on what is known at fixed points in time, and don't automatically adapt to change.

Lexicons can't 'learn' over time

Following on above, unlike **AI-based** surveillance systems, lexicon-based surveillance systems can't self-learn over time. Lexicon rules need to be manually curated and updated on a regular basis.

Lexicons aren't holistic

Because they only focus on communications surveillance, lexicon-based surveillance solutions don't have any inherent ability to include holistic surveillance data, or link to other sources of data.

4. HOW CAN NLP HELP?

Natural Language Processing (NLP) is a field of **Artificial Intelligence (AI)** that focuses on enabling computers to understand, interpret, and generate human language in a way that is meaningful and contextually relevant.

Alan Turing, a British mathematician and computer scientist, is often credited with laying the theoretical foundation for **NLP**. In his 1950 paper *Computing Machinery and Intelligence*, he introduced the concept of the Turing Test, a benchmark for evaluating a machine's ability to exhibit intelligent behavior, including natural language understanding.

Today, **NLP** is used to perform numerous tasks, ranging from basic (e.g. character recognition and speech recognition), to complex (e.g. text summarization and machine translation). The current widely discussed and debated concept of "**generative AI**" ("**GenAI**") refers to more advanced deep learning and the ability to not just understand existing content, but additionally generate new content based on understanding (including text, images, and other media).

How does NLP work in terms of alerts in financial services compliance?

NLP is reshaping the world of compliance and surveillance in financial services by helping firms analyze growing volumes of communications, across every communication modality, so firms can more accurately flag risky communications.

For example, NICE Actimize's **SURVEIL-X**[3] can understand and analyze communications in 45 different languages. It automatically detects people, places, products, companies, trades, assets classes and conversation topics within eComms and other electronic communications, as well as transcribed voice conversations, providing unique insight into what regulated employees said and did.

NICE Actimize's Director of Compliance Product Management, *Steven LoGalbo* explains: "When it comes to communications, conduct surveillance is especially challenging because it's unstructured content. **NLP** and machine learning are continuously evolving and getting more advanced and can help firms learn a lot about employees based on their communications. With **NLP**, we can see how people are interacting with one another, who they're interacting with and how frequently, what they're saying, and how they're saying it, by extracting context and sentiment from communications. This, in turn, can surface potential issues and provide an early warning to supervisors when behaviors aren't fitting normal patterns."

Fine-tuned for financial markets, **SURVEIL-X**'s **NLP** can even detect jargon indicative of inappropriate sales practices or aggressive behavior.

NLP can identify and alert to: negative sentiment; signs of fear, aggression and short-temperedness; pushiness on the part of a financial advisor; mistrust; client confusion and hesitation when discussing products; and even overuse or misuse of technical jargon which can be hard for customers to understand.

NLP can also analyze speech to understand the tone and intent of conversations, and to extract meta-data, such as quote or trade details from communications, so communications (unstructured data) can be more accurately correlated with trades (structured data).

3 More info: https://www.niceactimize.com/compliance/holistic-surveillance.html

Detection is the core of any surveillance system. Here are examples of some of the different risky behaviors that NLP can detect:

- Conversations containing promises, false assurances and guarantees
- Sludge tactics, like pressure selling, or intentional friction in communications
- Financial representative aggression, pushiness or client confusion and hesitation when discussing products
- Overuse or misuse of technical jargon, which can be hard for customers to understand
- Disclosures, both verbally and in written form, and whether they were timely and adequate
- Signs of aggression, confusion or hesitation when discussing products
- Communications indicative of market manipulation and/or collusive behavior
- Inappropriate and malicious internal communications (e.g. conduct policy violations)
- Improper sharing of information or data

NLP can also help to distinguish between material and non-material events. For example, **SURVEIL-X** can ingest data from over 19,000 news sources. Using **NLP** and entity extraction techniques, **SURVEIL-X** can identify each news's relevance, sentiment and potential market impact, and assign a relevancy score. The higher the score, the greater the likelihood it represents a material event.

Going one step further, **SURVEIL-X** can even use **AI** and **Advanced Analytics** to correlate employees actions (trades and behavioral data) with communications to help firms understand what employees said, heard and did, and uncover hidden risks more accurately and efficiently than ever before.

By leveraging **AI**-driven holistic surveillance solutions, firms can analyze communications, alongside trade and behavioral data, to detect and prevent market abuse and conduct risk, reconstruct events, and better understand the intent behind employee actions.

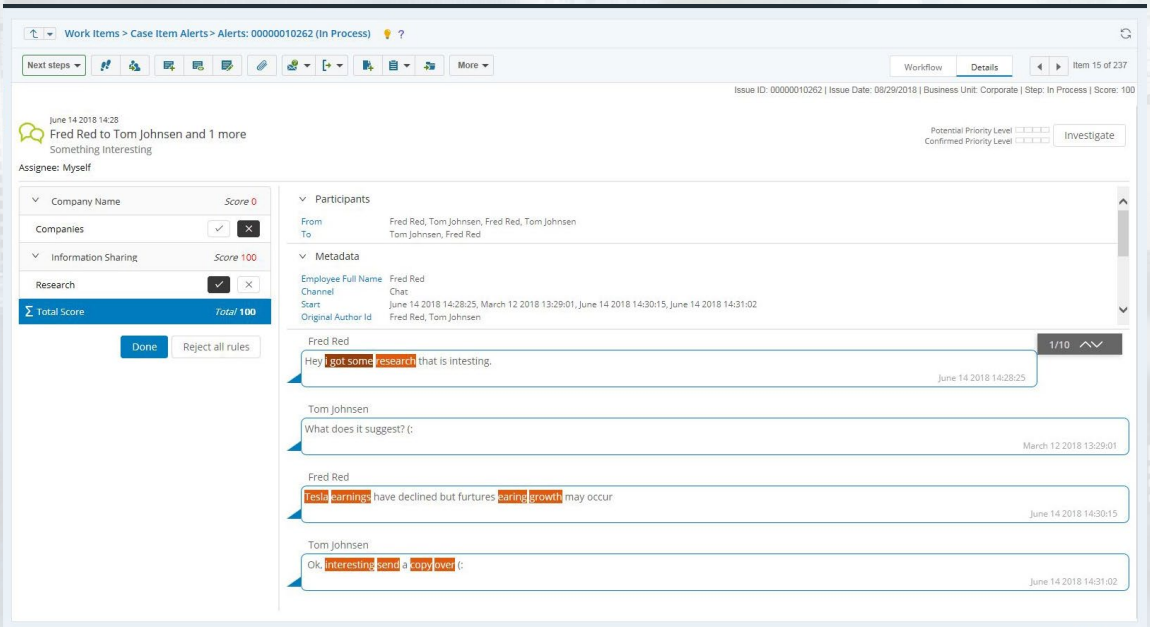

Screenshot of **SURVEIL-X** *Communications Surveillance with* **NLP**

5. WHY A HYBRID APPROACH IS BETTER

While most compliance teams still rely on lexicons exclusively, some are coming to the realization that using lexicons along with **Natural Language Processing (NLP)** is a better, smarter approach.

Why? **NLP** can assist organizations in overcoming all of the challenges of lexicons noted before.

But that's not to say that **NLP** should completely replace lexicons today. There's still a role for lexicons in communications surveillance. The beauty of lexicons is in their simplicity and ability to spot key words. Think 'Game Stop,' for example.

Lexicons are also viewed as a sufficiently proven methodology by regulators, a technology that is intuitive, and easy to explain.

And it's worth mentioning, that while newer technologies like **NLP** have been shown to reduce false positive alerts (over lexicons), the very act of producing fewer alerts can be construed by regulators as potentially missing issues.

Still, some regulators are now starting to demand that firms update their methodologies to incorporate more modern approaches to identifying risks, such as **NLP** and **Machine Learning (ML)** technology.

A hybrid approach makes sense for other reasons too. Lexicons are fundamental to developing robust **NLP** models. Lexicons play a very important role in "training" and "seeding" new **NLP** and **AI/ML**-based models with baseline information, enabling models to have a foundational understanding of language and industry jargon, and enhancing their accuracy.

"**NLP** models need labeled data to flag risky communications; lexicons are the starting point in labeling the data," explains *Nitin Vats*, Product Manager – Data & AI, NICE Actimize."

The fact is, many firms already have historical labeled data based on the lexicon rules. So, for example, if one of the lexicon rules is to flag communications based on keywords like "insider trading," "personal trade," etc., these flagged interactions can be the starting point to train the **NLP** model.

For example, Vats says, firms can create a dictionary of lexicons that can be used to label and train the **NLP** model, including:

- Domain-specific lexicons, which can be used to label text data that is unique to a specific financial services industry or context
- Offensive or profane words that can be used to flag and label text data containing inappropriate content
- Words that can be used to identify and extract entities (e.g. people, organizations, locations, etc.)

NLP models can also be trained using lexicons, along with the user reviewed data. For example, a compliance analyst might manually label 100 examples of aggressive or secretive interactions (identified via lexicons) and then feed them into the **NLP** model. Manually reviewed interactions are generally more accurate then interactions identified by lexicons (that are never reviewed by a human).

But, Vats cautions firms to make sure that any flagged interactions (whether based on lexicons alone, or on user reviewed data) are thoroughly reviewed for accuracy before using them to train any **NLP** models.

One other advantage of using **NLP**-based surveillance is the ability to perform sentiment analysis on communications, but this requires the use of lexicons too. Absent lexicons, one can't attribute specific sentiment to words. Lexicons enable some of the more advanced capabilities of **NLP**.

While lexicons can serve as a foundation to **NLP**, **NLP** can also work in a supporting role for lexicon-based surveillance.

For example, with false positive alerts remaining one of the biggest challenges for firms, **NLP** models can be used to automatically sift through an initial set of alerts generated by an existing lexicon-based surveillance tool. This helps to eliminate false positives, and produce a more manageable, accurate set of alerts for investigation and review.

In this manner, firms are essentially using lexicons as a first line of defense, and leveraging **AI**/**NLP** as a filter to weed out false positives.

Firms also might want to incorporate additional granular rules, such as...

> -*Convert an* **NLP**- *flagged risky interaction into an alert when the sentiment is negative with a score of 90 and above, and the intention is secretive or aggressive*

All of this aside, the broader concept of leveraging lexicons and **AI** together is providing hope for firms, especially with respect to monitoring eComms.

As the ***2023 Solving Surveillance report*** from **1LOD** explains: "In e-comms, **AI** solutions can be used to process the vast data sets generated and identify the higher-risk alerts, potentially paving the way for a more sensible, risk-based approach to surveillance and replacing the current 'review everything' model."

Hybrid Approach to Comms Surveillance Yields Fewer False Alerts and Better Results

In theory, using lexicons and **NLP** together sounds like a good idea, but does it produce results?

Actual results from a sell-side bank confirm that using lexicons and **NLP** together can dramatically reduce the number of alerts that need to be reviewed and investigated, while markedly improving a compliance organization's ability to identify true risks (see chart below).

Using lexicon-based alert generation plus **NLP**, one sell-side bank was able to reduce false positives by a remarkable 72%, and reduce the number of daily alerts that needed to be reviewed from one thousand to ten. In addition to reducing false alerts, the bank was able to identify three new risks that it hadn't previously known about. And the compliance team was able to conduct investigations four times faster.

Lexicon-based alert generation	Lexicon-based alert generation plus NLP
• 1,000 alerts generated per day • 99.999% of alerts generated are false positives	• 280 alerts per day (72% false positives reduction) • 270 alerts predicted to be false positives • 10 alerts requiring review • 3 new true risks identified • 4x faster investigations

Like all shiny new things, at first glance, the impressive capabilities of **NLP** can make other detection solutions appear instantly obsolete. But whether **NLP** is on your short or longer-term surveillance trajectory, there's a lot of value that can be achieved by leveraging lexicons and **NLP** together.

Adoption of NLP: Proof in Numbers

According to a recent NICE Actimize webinar survey, 25% of respondents are already actively using **NLP** in their communications surveillance today.

In a separate **Chartis Research Survey**[4]: *The Future of Trader Surveillance*, financial services firms are now prioritizing **AI** (machine learning and behavioral analytics) as essential to trade surveillance. In fact, close to 50% of firms surveyed listed **AI** as a key technology driver. While this study focused on trade surveillance, there is an undeniable similarity to **NLP** applications for comms surveillance.

What priority do you expect the following to have over the next 3 years in order to reduce the manual effort to review false positives ? (High, Medium, Low)

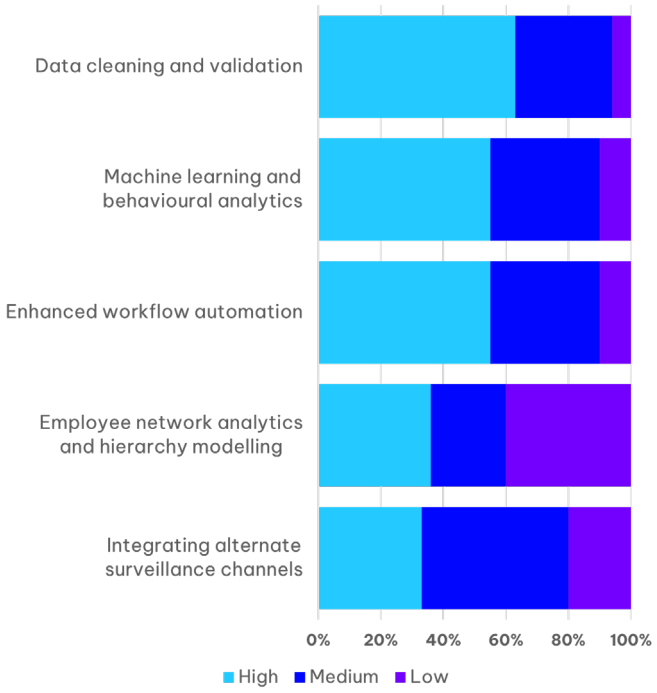

Chart: from Chartis Research Survey

According to a separate **Opimas report**[5], *Definitive Guide to Modern Trade and Communications Compliance*, "Very few firms are satisfied with the quality of the alerts they produce. Firms striving for the best possible surveillance results, and a manageable number of alerts, should carefully combine classic lexicons, meta-data, trade and order information, as well as more cutting-edge techniques that enable language identification, entity recognition, sentiment analysis, noise deduping, topic identification, clustering, and more."

The Opimas report goes on to say, "Supervised and unsupervised machine learning should be applied to better categorize communications and trading activity that warrant attention, helping to save time. These techniques can also be used to implement a system of scoring alerts for riskiness, which can further help to prioritize investigations for analysts."

Deloitte's *Market Abuse Outlook 2022* also affirms that regulators around the world are also now embracing **Machine Learning (ML)** and **AI** to identify patterns, trends and anomalies.

4 https://www.chartis-research.com/regulatory-compliance-and-reporting/financial-reporting/future-trader-surveillance-1189

5 https://actimize.nice.com/opimas-compliance-guide.html

Benefits of Using NLP Along with Lexicons

Aside from the results mentioned earlier in this White Paper, the combination of employing **NLP** along with lexicon-based approaches to communications surveillance affords many other advantages:

Gives structure to unstructured data

NLP models can organize and structure data in documents, voice transcripts, chat and other unstructured data sources for easier review and processing.

Reducing noise

NLP models are proven to reduce noise by up to 90%, and identifying and removing noise is critical to reducing false alerts. For example, **NLP** models can be trained to recognize non-relevant content and flag it for removal, or automatically remove it from view.

For example, **NLP** disclaimer removal models are trained on historical disclaimer content, and as such, can easily identify new disclaimer messages (without the need to include various disclaimer language in a lexicon dictionary). Once a disclaimer message is identified, it can be removed.

Additionally, using a lexicon-based surveillance approach, every time an email is forwarded it can end up creating new alerts. **NLP**-based duplicate identifier models can easily identify whether a particular email is a duplicate, and what percent of the content is new or duplicated. When true duplicate emails are identified, they can also be removed, thus reducing the potential for additional false alerts.

Understanding context of communications, identifying true risk

Accurate alert prediction relies on understanding the context of communications. **NLP** overcomes lexicon's limitation of not being able to understand the context of communications.

Understanding the context of communications is a crucial aspect of **NLP**, and it involves several techniques and processes, including part-of-speech tagging (the ability to recognize roles of words in sentences), dependency parsing (the ability to understand grammatical relationships between words in a sentence), all coupled with the ability to recognize entities and analyze sentiment. Machine learning on deep data sets also is crucial to an accurate understanding of context.

Identifying new risks

Communications are dynamic in nature. Language conventions change. The devices and communication channels regulated employees use change over time as well. Employees can also find innovative ways to bypass surveillance (e.g. switching from a voice call to chat, or using a different language). With all of these changes, it can be tough to identify new risk indicators.

By identifying anomalies, recurring topics, and emerging patterns in communications, **NLP** can accurately detect new risks before they can adversely affect your firm.

Detecting entities, sentiment, and intent

Understanding what is being said, and the emotion and intent behind what is being said, is essential to effective communications surveillance.

This is an area where **NLP** excels.

NLP uses a technique called entity extraction to identify and classify named entities found in communications, into predefined categories such as specific people (who's speaking or being spoken to, or about), company names, locations, dates, monetary values, and so on.

After all, it's difficult to flag true risks, if you don't know what's really being said.

Scientific research indicates that the presence or absence of specific types of sentiments expressed in communications can be an indicator of increased potential for conduct risk too. **NLP** can detect and classify communications containing a range of sentiments and emotions – fear, unfairness, mistrust and aggression – that can be early indicators of conduct risk. Using **NLP**, the surveillance system can assign a confidence score to each detected sentiment or emotion.

Using these techniques to understand the context of communications can also help to reveal the intent behind interactions.

Understanding trader jargon

Because they can be trained on financial data (communications) across many languages, **NLP** models can be fine-tuned to understand nuances in language that are specific to the financial services sector, including trader jargon, and abbreviations. For example, the term "Kiwi" may refer to a fruit, but it's also used to refer to the New Zealand dollar.

Increasing investigative efficiency

In addition to reducing false alerts, **NLP**-powered surveillance can assign degrees of confidence to alerts, which can significantly cut down on compliance analysts' workloads by allowing them to focus their investigations on alerts that have a higher probability of being true. This strategy certainly pays off. In one study, a financial services firm was able to save 52% on their investigation and alert review cost by leveraging **NLP**.

Deloitte's *Market Abuse Outlook 2022* highlights the benefits of this approach as well. The report states: "**AI**-based systems provide risk scores, allowing for improved incident prioritization and categorization. This could significantly reduce false positives and improve the effectiveness of alert investigations. These efforts are not just limited to trade surveillance. Communication surveillance capabilities have also been upgraded with **AI** techniques such as **Natural Language Processing** (**NLP**). Historical alerts are analyzed and categorized using **NLP** to create a profile of what a high-quality alert looks like. Following that, **Machine Learning (ML)** models are used to adjust and optimize surveillance parameters to maximize surveillance output, with the goal of improving the overall quality of the surveillance alert pool and reducing false positives."

Improving accuracy over time with machine learning

Additionally, with machine learning **NLP** models get smarter over time, by using analyst feedback for previously reviewed alerts to improve future detection accuracy. Machine learning techniques are a fundamental component of **NLP** models, and contribute significantly to their accuracy and effectiveness.

Streamlining event reconstruction

NLP also enables faster, automated correlation of communications and market surveillance data for speedy trade reconstruction, so it's easy to verify what people said and did.

6. TIPS FOR SUCCESSFULLY DEPLOYING A HYBRID APPROACH

Ready to embrace a smarter communications surveillance approach? Here are some tips to ensure your success, collated into three key categories: *strategy*, *technology*, and *user experience*.

STRATEGY

The right strategy is what's right for your firm

When it comes to implementing lexicons and **NLP** together, the right strategy is what's right for your firm. For example, as mentioned above, you can leverage **NLP** models to automatically sift through an initial set of alerts generated by your existing lexicon-based surveillance tool to eliminate false positives, producing a more manageable, more accurate set of alerts to be investigated. Or you can begin with this strategy, with the aim of graduating to a communications surveillance program powered solely by **NLP**.

Leverage the experience of an outside company with expertise in surveillance and NLP

In addition to possessing institutional and regulatory knowledge, and purpose-built solutions for communications, behavioral, trade and holistic surveillance, a vendor can assist you with any data challenges and help you understand where and how to apply **NLP** and machine learning in a manner that can add value to your surveillance program.

Have a clear strategy

According to a **McKinsey & Company** survey, one of the most significant roadblocks to expansion of **AI** is the lack of a clear strategy. This is particularly true in the compliance arena where regulatory requirements and internal policies introduce complexities for surveillance and monitoring. When **AI** is implemented with a lack of understanding and strategic focus, this can lead to less-than-optimal results.

You can't solve a problem without defining it first, so start by defining the problem.

Natural language processing can be very effective in analyzing trader communications. That said, every surveillance challenge isn't necessarily best addressed through **AI**.

If you are going the **NLP** route, you'll also need to have uniform processes in place to ensure data quality. Data volume is an important starting point for supervised machine learning, but the biggest challenge really lies in the data quality. Machine learning can only make accurate predictions based on what it learns and knows in the first place. Bad data yields bad results. If your firm is considering deploying supervised machine learning, you should first make sure there are processes in place to ensure that data will be labeled uniformly and accurately. Prioritizing data quality also means instituting on-going review processes and data quality checks.

TECHNOLOGY CONSIDERATIONS

Buy in lieu of building

Deciding whether to buy a communications surveillance solution or build one from scratch depends on numerous factors, including your organization's specific needs, resources, and expertise. Both options have their pros and cons, and the choice should be made carefully based on your unique circumstances. There are many reasons why buying a communications surveillance solution may be a wiser choice than building one from scratch, including:

- **Cost-efficiency:** There's no need to hire skilled developers and engineers. Buying a ready-made solution also comes with a predictable upfront cost and lower total cost of ownership. On the other hand, building your own solution can lead to cost overruns and unpredictable expenses.
- **Time-to-market:** Building a surveillance solution takes time, from initial design and development to testing and deployment. Buying an existing solution can significantly reduce the time it takes to implement surveillance capabilities.
- **Expertise and support:** Surveillance solutions from reputable vendors are developed and maintained by teams of experts with extensive experience in the field.
- **Compliance and legal considerations:** Surveillance solutions often have to adhere to various legal and regulatory requirements, such as data privacy laws. Established vendors are more likely to have built compliance features into their products, saving you the effort of ensuring compliance on your own.
- **Scalability and updates:** As your needs evolve, it may be easier to scale a purchased solution up or down, and updates and improvements are often provided by the vendor.
- **Reduced risk:** Established solutions undergo rigorous testing and may have a better track record for security and reliability.
- **Integration:** Many surveillance solutions are designed to integrate seamlessly with other systems and data sources. Building a custom solution requires much more effort to achieve this level of integration.
- **Core competencies:** If surveillance is not your organization's core competency, buying a solution allows you to focus on what you do best while relying on experts for what they do best: surveillance tech.

Look for a solution trained on domain specific, financial data

NLP models trained on domain-specific data tend to perform significantly better within that specific domain compared to general-purpose models. They are more attuned to the specific vocabulary, jargon, and language patterns used within that domain. Domain-specific data also helps **NLP** models understand context and meaning more accurately. This is crucial for tasks like sentiment analysis, where context is key to sentiment.

Oftentimes firms will want to use their own data to train their **NLP** model, but in most cases, at least at the start of a surveillance deployment, that data is too scarce. For that reason, firms should look to source a solution from a firm with an extensive client base aligned with their domain. This federated strategy to modeling enables prediction models to train on different data sets (multiple clients/extensive industry data) as opposed to relying on models that are trained solely on limited individual client data sets.

The federated strategy benefits by having shared global models that leverage broader, deeper, diverse data sets.

Over time, the alert prediction model can be tuned to your firm's individual data, when that data becomes more abundant.

Make sure the NLP / AI models are explainable

Many view **AI** as a 'black box' of sorts, which data goes into and decisions magically come out of. You wouldn't blindly trust a machine to make consequential compliance decisions for your firm. That's why **AI** explainability is essential.

Fortunately, because alert prediction analyzes a myriad of data, predictions can be enriched with information, which ultimately helps compliance analysts save time and make better decisions. This enriched information enables compliance analysts to prioritize alerts based on their relevance and importance.

Look for a solution that incorporates drill down dashboards that provide detailed explanations of every alert prediction. The solution should provide a view into all the factors that contributed to the alert prediction, (for example, conversation language, conversation topics, and so on, along with each factor's relative weight).

Knowing what went into each alert score can reduce compliance analyst workload and help them make better decisions. At the end of the day, it can also give your firm greater confidence in **AI**, and go a long way toward refining, improving and retraining your **NLP** models for better accuracy.

Look for a solution with self-development capabilities

As communications evolve over time, so do risk profiles. Having self-development capabilities allows **NLP** systems to adapt to changing requirements, stay relevant, and maintain their usefulness over time. Therefore firms might want to look for solutions that also include easy-to-use self-development toolkits to make training, building, and deploying **NLP** models fast and easy.

Using self-development tools, firms can:

- Create new models and retune existing ones quickly to address their firm's specific business needs
- Test and validate **NLP** models before full-scale implementation
- Reduce the high costs associated with building **NLP** models (which are typically dependent on specialized **AI** skill sets)
- Process and analyze data instantaneously for informed decision-making

Look for a solution that can provide for a smooth migration path to holistic surveillance

Surveillance and decision-making is all too often done in silos. This fragmentation of controls means there's no overall big picture view of regulated employee actions which would enable firms to accurately identify issues. For example, different systems may set off alarm bells when a trader breaches his volume limits, has unusual patterns of cancellations or corrections, exhibits other deviations from normal trading patterns, or makes unauthorized P&L adjustments. But traders also talk (text or chat) about the trades they're doing, too, and these communication alerts are captured in other systems. There's no way to automatically link all of this data together.

With a holistic strategy you can eliminate siloed decision-making by leveraging **AI**-powered predictive algorithms to analyze behavioral, trading and communications data, and detect and deter conduct issues before they get out of hand.

USER EXPERIENCE

Change management is important when introducing NLP

Make sure people understand their roles and what to expect.

Change management is critically important in implementing **Natural Language Processing (NLP)** for several reasons:

- **Organizational transformation:** Implementing **NLP** often involves significant changes in how an organization operates. It can affect processes, workflows, and roles within the organization. Change management helps guide this transformation and ensures that everyone is on board with the changes.
- **User adoption: NLP** systems are only effective if they are used. Change management helps to address resistance to change, encourage user adoption, and make the transition smoother for everyone involved.
- **Cultural shift: NLP** can introduce a cultural shift in an organization, especially if it automates tasks that were previously done manually. Change management helps in communicating the reasons behind the shift, aligning the culture with new goals, and ensuring that employees are comfortable with the changes.
- **Training and skill development: NLP** implementation often requires employees to acquire new skills and knowledge. Change management can facilitate training programs, workshops, and resources to ensure that employees have the necessary competencies.
- **Risk mitigation:** Change can introduce risks, such as disruptions in operations, loss of productivity, or data security concerns. Change management strategies can identify these risks early and provide plans to mitigate them.
- **Stakeholder communication:** Effective change management involves clear and transparent communication with all stakeholders. It ensures that everyone understands the goals, benefits, and impacts of **NLP** implementation.

- **Feedback loop: NLP** relies on supervised machine learning, and humans (namely compliance analysts) play an essential role in providing feedback on alerts, which in turn improves alerting accuracy. Alerts need to be properly dispositioned. If they're not, the surveillance system will make future predictions based on bad data.
- **Change sustainability:** Change management doesn't end once the **NLP** system is implemented. It also involves ongoing monitoring and support to ensure that the changes are sustainable in the long term.
- **Compliance and ethical considerations: NLP** systems often deal with sensitive data, and their use may be subject to regulatory requirements and ethical considerations. Change management can help ensure that the implementation aligns with these requirements and values.
- **Measuring success:** Change management establishes key performance indicators (KPIs) and metrics to measure the success of the **NLP** implementation. This allows the organization to evaluate the return on investment and make necessary adjustments.

Prioritize user experience

At the end of the day, human beings who are responsible for implementing the technology need to understand what's working and what isn't. Predictions made by unsupervised models can and should be reviewed regularly to ensure that the models are in fact making accurate predictions. For this reason, if you're considering deploying **NLP** and machine learning for communications surveillance, look for a solution that incorporates real-time dashboards that provide a seamless experience for your staff, and allow them to easily monitor your supervised machine learning models' accuracy.

This will allow them to see which models are accurate (based on dispositions of alerts) and which are not, or how a particular model is performing compared to others. Then they can use this information to make necessary tweaks and adjustments.

You'll also want to make sure that user experience is seamless for your compliance analysts, too, so look for a solution that distills alert data into color coded dashboards. Using these dashboards, analysts should be able to filter employees based on scores, and drill down into employee data to view trend timelines, behavioral spikes and contributing factors. The drill-down dashboard should also give them insight into underlying communications.

Double down on model risk management capabilities

Risk management and **Artificial Intelligence (AI)** are closely intertwined in today's financial services landscape.

Model risk management has traditionally been used to manage financial and credit risk models but has now expanded into areas where **AI** and **ML** models are used. Regulators expect companies to be able to understand and explain exactly what their models do. Company leaders have similar expectations as well. The model risk management function oversees these things.

For this reason, Accenture's Frieder suggests doubling-down on model risk management capabilities to make sure that you are not developing something that's unmanageable.

Most financial services firms already have model risk management functions in place to identify, assess, monitor, and mitigate potential risks associated with the use of various models, including statistical, financial, and **AI** models. Model risk management performs the critical role of testing the accuracy and reliability of models, ensuring they are appropriate for their intended purposes, and maintaining ongoing oversight to prevent adverse outcomes based on flawed models.

Follow best practices for AI ethics and governance

Ethics and governance in Natural Language Processing requires firms to address the unique ethical considerations and regulatory challenges that arise when dealing with **AI** systems that are used to understand, generate, and manipulate human language. Here are some best practices to ensure your responsible use of **NLP**:

- **Consider the ethical implications before developing and deploying NLP models and algorithms.** For example, could your system be used to discriminate against certain individuals or groups? Could it be used to spread misinformation or propaganda?
- **Use fair and unbiased data.** The data that you use to train your **NLP** system can have a substantial impact on its performance. If your data is biased, your system is likely to learn to be biased as well. To avoid this, make sure that your data is fair and unbiased. This means that it should represent a diverse range of viewpoints and experiences.
- **Be transparent about how your model or algorithm works.** It is important to be transparent about how your **NLP** system works. This means providing information about the data used to train your system, the algorithms that you used, and the limitations of your system. Document the data sources, preprocessing steps, and algorithms used to develop your **NLP** models. Transparent documentation helps ensure accountability and makes it easier for others to understand and review your work. By being transparent, you can help users understand how your system works and to make informed decisions about how to use it.
- **Respect the privacy of individuals.** When you are using **NLP** to collect or process personal data, it is important to respect the privacy of individuals. This means obtaining consent from individuals before collecting their data and taking steps to protect their data from unauthorized access or disclosure. Follow strict data privacy standards, especially if your **AI/NLP** system deals with personal or sensitive data. Implement encryption, access controls, and other security measures to protect user data.
- **Hold yourself accountable.** As a developer or user of **NLP**, you have a responsibility to use these technologies in a responsible and ethical manner. If you become aware of any potential harms that your system could cause, take steps to mitigate those harms. Clearly define roles and responsibilities within your team for ensuring the responsible development and deployment of **NLP** models. Hold individuals and teams accountable for ethical lapses. By holding yourself accountable, you can help to ensure that **NLP** is used for good.
- **Provide education and training on NLP and AI/ML.** Provide training to your team on ethical considerations in **NLP** and **AI/ML** development to foster a culture of ethical awareness and responsibility.
- **Maintain human oversight and intervention** in the use of **NLP** and **AI/ML**, especially in critical decision-making processes. Human experts should be available to review and intervene in complex or sensitive situations.
- **Regularly monitor and maintain/update your NLP and AI/ML models and algorithms.** Continuously monitor the performance of your **NLP** models in real-world scenarios. Update and retrain models as needed to maintain their accuracy and effectiveness over time.
- **Provide "informed consent"** if your **NLP** model is used to interact with users. Make sure users are informed about how their data will be used and obtain their consent. Provide clear explanations about the purpose, scope, and potential implications of using the technology.
- **Perform regular "ethical reviews."** Establish an ethics review process for **NLP** and **AI/ML** initiatives, especially those that have the potential to impact society or customers. Consider forming ethics committees or seeking external input to evaluate the ethical implications of your work.
- **Stay current with relevant NLP and AI/ML related regulations** and standards in your industry and region. Ensure that your systems comply with these regulations.

Take a proactive stance in addressing bias in NLP models

Addressing bias in Natural Language Processing models is a critical and ongoing challenge in the field of artificial intelligence. Bias in **NLP** models can manifest in various ways, including gender, race, religion, and cultural biases. Mitigating these biases is essential to ensure that **AI** systems are fair, ethical, and equitable. Here are some of the key strategies to address bias in **NLP** models:

- **Diverse and representative data:** Ensure that the training data used to build **NLP** models is diverse and representative of the real-world population. Biases in **NLP** models often result from biased training data.
- **Bias detection in data:** Employ bias detection techniques to identify and quantify biases in the training data. Tools like fairness audits and bias detection algorithms can help in this process.
- **Debiasing data:** Use techniques such as re-sampling, re-weighting, or data augmentation to reduce bias in training data. For example, you can balance the data set to ensure equal representation of different groups.
- **Fairness constraints:** Integrate fairness constraints into the training process. This can involve adjusting the loss function to penalize predictions that exhibit bias or fairness violations.
- **Adversarial training:** Implement adversarial training, where an additional component is added to the model to counteract biases in the main model.
- **External audits:** Conduct external audits of **NLP** models by independent parties to assess their fairness and bias.
- **Ethical guidelines and frameworks:** Adhere to ethical guidelines and frameworks when developing **NLP** models, such as the ACM's Code of Ethics and Professional Conduct or the IEEE's Ethically Aligned Design.
- **Transparency and explainability:** Ensure that **NLP** models are transparent and explainable. Users should be able to understand how a model arrived at its decisions and identify any biases in the process.
- **Continuous monitoring and iteration:** Continuously monitor and re-evaluate models in real-world applications to identify and address any emerging biases.

Stay informed on regulations governing the use of AI/NLP

The regulations governing the use of Natural Language Processing (**NLP**) technology can vary significantly from one jurisdiction to another and are subject to change over time. Here are a few things you need to be aware of:

- **Privacy: NLP** models and algorithms can be used to collect and analyze large amounts of personal data, which can raise privacy concerns. Again, regulations in this area typically require companies to obtain consent from individuals before collecting their data, and to take steps to protect the data from unauthorized access or disclosure.
- **Discrimination and bias: NLP** models and algorithms can be used to make decisions that may potentially be viewed as discriminatory, such as decisions about hiring, lending, or insurance. It is important to be aware that **NLP** models can inherit biases from the data they are trained on. As noted above, firms should conduct regular audit and evaluation of models for bias and fairness.
- **Security: NLP** models and algorithms can be used to access and control sensitive systems and data, which can raise security concerns. Regulations in this area typically require organizations to implement appropriate security measures to protect their systems and data from unauthorized access or attack.
- **Transparency: NLP** models and algorithms can be complex and opaque, making it difficult for individuals to understand how they work and how data is being used. In general, regulators expect companies to have in-depth understanding of what their models do and how they use data, and to be able to verify they are working as intended.
- **Employment and labor laws:** If **NLP** is used in the context of hiring, employee evaluations, or workplace communications, there might be implications related to employment and labor laws, particularly those concerning fairness, discrimination, and privacy.

7. MAKE THE MOVE TO SMARTER SURVEILLANCE

This White Paper covered the considerations around **Artifical Intelligence (AI)** and in particular **Natural Language Processing (NLP)** within surveillance, and some best practices and tips to get you started. But as you embark on your **AI** journey toward smarter communications surveillance, challenges will come up.

As the largest and broadest provider of financial crime, risk, and compliance solutions for regional and global financial institutions, ***NICE Actimize*** can help. The supervised and unsupervised machine learning solutions from NICE Actimize have been successfully deployed at many leading financial institutions. No compliance technology vendor is better equipped to help you understand where and how to apply **AI** and **Machine Learning (ML)** for optimal surveillance results. Finally, NICE Actimize's surveillance drill-down dashboards remove the mystery of **AI** by providing complete explainability and confidence scores for every alert.

Accenture can also help. As a renowned leader in guiding companies on their **AI** and **ML** journey, and with a strong reputation for providing comprehensive consulting services and innovative solutions, Accenture has a track record of successful implementations across various industries. Their expertise and thought leadership in **AI** and **ML** make them a top choice for organizations seeking support in adopting these technologies.

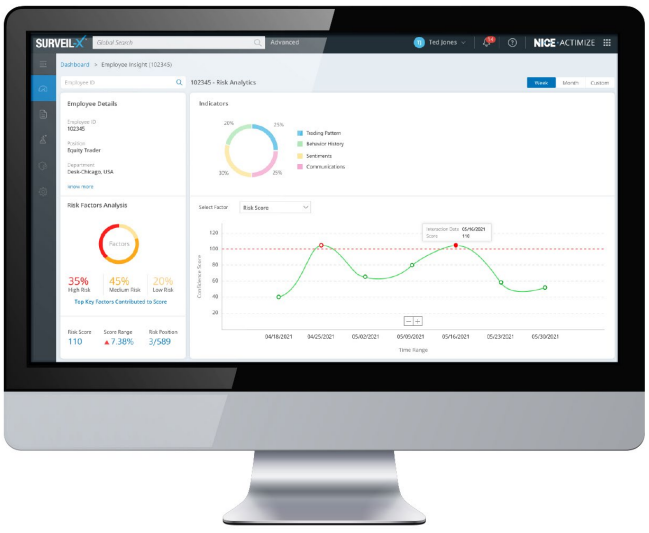

Contact the contributors to this White Paper

Nitin Vats

Product Manager,
Data & AI

NICE Actimize

Nitin.Vats@nice.com

Jonathan Frieder

Principal Director,
U.S. Regulatory and Compliance Practice

Accenture

Jonathan.Frieder@accenture.com

About Accenture

Accenture is a leading global professional services company that helps the world's leading businesses, governments and other organizations build their digital core, optimize their operations, accelerate revenue growth and enhance citizen services—creating tangible value at speed and scale. We are a talent and innovation led company with 732,000 people serving clients in more than 120 countries. Technology is at the core of change today, and we are one of the world's leaders in helping drive that change, with strong ecosystem relationships. We combine our strength in technology with unmatched industry experience, functional expertise and global delivery capability. We are uniquely able to deliver tangible outcomes because of our broad range of services, solutions and assets across Strategy & Consulting, Technology, Operations, Industry X and Accenture Song. These capabilities, together with our culture of shared success and commitment to creating 360° value, enable us to help our clients succeed and build trusted, lasting relationships. We measure our success by the 360° value we create for our clients, each other, our shareholders, partners and communities. **www.accenture.com**

About NICE Actimize

NICE Actimize is the largest and broadest provider of financial crime, risk and compliance solutions for regional and global financial institutions, as well as government regulators. Consistently ranked as number one in the space, NICE Actimize experts apply innovative technology to protect institutions and safeguard consumers' and investors' assets by identifying financial crime, preventing fraud and providing regulatory compliance. The company provides real-time, cross-channel fraud prevention, anti-money laundering detection, and trading surveillance solutions that address such concerns as payment fraud, cybercrime, sanctions monitoring, market abuse, customer due diligence and insider trading.

niceactimize.com/compliance | compliance@niceactimize.com