WHITE PAPER

Using Analytics to Proactively Deter Insider Threats

Strengthen organizational values in times of high security risks

sas
THE POWER TO KNOW.

Contents

Stories about security challenges in both government agencies and commercial businesses have become a mainstay in today's media. Hardly a day goes by without a major announcement of a data breach, stolen personally identifiable information (PII) or compromised records. While many of these malicious acts are perpetrated by hackers and fraudsters in order to make a political statement or for financial gain, an increasing number are carried out by insiders – trusted employees, contractors or business partners – for many of the same reasons.

Hundreds of software vendors offer solutions that cater to these types of threats, making the determination of an effective defense-in-depth strategy investment seem especially overwhelming. However, there are risk-mitigating strategies that your organization can implement to complement your existing security, ethics and compliance, and internal audit programs.

Unlike its external equivalent in the cyber domain, internal threats have posed risks to organizations since the first employee was hired. Advances in technologies, the monetization of sensitive data and common stressors in human nature all have contributed to the rise in insider losses. For decades, this type of threat seemed unique to the national security market. Consider such famous cases as Robert Hanssen at the FBI, the self-recruited traitor who was the focal point in the 2007 movie *Breach*, and the self-proclaimed whistleblower Edward Snowden, an NSA systems administrator contractor who stole massive amounts of data detailing highly secretive technical collection programs. Although corporate espionage does not typically make front-page news, many examples exist where competition takes precedence over ethical or legal obligations, including:

- August 2016: Ongoing legal proceedings where Jawbone accused Fitbit of patent infringement and hiring key employees away.[1]
- January 2015: Morgan Stanley Wealth Management employee Galen Marsh is fired over stolen (or mishandled) client data.[2]
- July 2014: Jun Xie, a Chinese engineer working for a subsidiary of GE Healthcare, stole 2.4 million files on trade secrets and other sensitive company data, which the FBI and GE alleged were sent to China.[3]

1 Dalton, Andrew, "US Judge absolves Fitbit of corporate espionage allegations (for now)," Engadget, Aug. 24, 2016. engadget.com/2016/08/24/us-judge-absolves-fitbit-of-corporate-espionage-allegations/

2 Baer, Justin, "Morgan Stanley fires employee over client-data leak," *The Wall Street Journal*, Jan. 5, 2015. wsj.com/articles/morgan-stanley-terminates-employee-for-stealing-client-data-1420474557

3 Vielmetti, Bruce, "Chinese engineer accused of stealing trade secrets from GE unit," *Milwaukee-Wisconsin Journal Sentinel*. archive.jsonline.com/news/crime/chinese-engineer-accused-of-stealing-trade-secrets-from-ge-unit-b99344912z1-274122821.html

- June 2010: Dyson employee Yong Pang allegedly sold competitive information to German rival Bosch.[4]
- 2006-2007: HSBC computer technician and whistleblower Hervé Falciani downloaded the details of 130,000 holders of secret Swiss accounts. This information was turned over to French investigators in December 2008 and then circulated to other European governments. It was eventually used to prosecute tax evaders.[5]

As explored in this paper, investing in the right analytical capabilities to detect suspicious behavior is far less costly than the damage of a successful insider threat.

Evolve and Expand Insider Threat Programs

The majority of safeguards for insider threat focus on detecting data loss, which is a risky approach. This is because technology enables data to be transferred in near-real time, and post-event mitigation is very costly.

In the case of insider threats, where employees, contractors and even business partners are subject to policies, procedures and controls, employers can apply traditional security approaches that deter bad actions or poor judgement by the employee. This approach reduces organizational risk while helping the individual remain lawful and trustworthy – a win for both parties.

While insider threat detection has been integrated into many sensitive programs, most organizations still handle insider threat infractions in an ad hoc, reactive manner. Even with the most mature and robust programs that incorporate software, processes are often still manually driven and inefficient. In addition, they focus on network or host activity (such as rule violations) with little consideration for human behavioral issues.[6] The systems used are fragmented and limited in value because they are:

- **Based on rules and driven by policy deviation** – detecting predefined events, such as numerous unsuccessful login attempts.
- **Focused on reactive investigations** – relying upon tips or leads, usually responding after an event occurs.
- **Focused only on network or host activity** – for example, for monitoring websites visited and files downloaded. Nontechnical information, such as badging records or phone records, often are reviewed separately, preventing a more holistic view.
- **Lacking prioritization** – which can be problematic with a large number of alerts overwhelming investigators, making it difficult to identify the most severe threats.

4 Gardham, Duncan, "'Spy' at the centre of Dyson espionage case named," *The Telegraph*, Sept. 27, 2016. telegraph.co.uk/finance/newsbysector/retailandconsumer/9701583/Spy-at-the-centre-of-Dyson-espionage-case-named.html

5 Garside, Juliette, "HSBC whistleblower given five years' jail over biggest leak in banking history," *The Guardian*, Nov. 27, 2015. theguardian.com/news/2015/nov/27/hsbc-whistleblower-jailed-five-years-herve-falciani

6 *A Preliminary Examination of Insider Threat Programs in the US Private Sector*, Intelligence and National Security Alliance – Cyber Council: Insider Threat Task Force, September 2013.

The Solution: Proactive Deterrence Powered by Hybrid Advanced Analytics

In order to remain vigilant – or at a minimum, keep pace with adversaries – organizations need to take a multifaceted, automated approach to threat evaluation that combines robust data management capabilities, a hybrid analytical approach, and an automated environment.

The fact is, no single analytics approach detects threats effectively. Motivations and tactics continually evolve, making the creation and maintenance of rules-based systems time-consuming and expensive. SAS® incorporates multiple methods – business rules, anomaly detection, predictive analytics and network analysis – so that unusual and suspicious user behavior can be surfaced for proactive deterrence and prioritized investigations. This approach also reduces the amount of false positives, enabling investigators to work many times the number of cases and focus on higher-risk networks. Additionally, SAS provides a feedback loop via decision outputs that are fed back into the analytical models to ensure they remain effective and relevant over time.

Beyond these capabilities, SAS offers a new technology platform, SAS® Viya™, allowing data to be accessed on a server, a grid or the cloud, and a new end-user investigative interface, SAS Visual Investigator. These offerings provide the most streamlined data analysis process combined with high-performance analytics for the fastest and most reliable insights.

The combination of powerful data aggregation, a hybrid analytical approach, and a powerful technology platform enables organizations to assume a more proactive and contextually aware security posture. This approach is critical to circumventing a potential terrorist plot, thwarting an espionage mission, reducing fraud transactions, and addressing other complex threats.

SAS: A Force Multiplier

SAS delivers unique, end-to-end threat detection based on best practices and intellectual property. It uses powerful data management and analytics technologies that are fully integrated with automated processes, enabling accelerated, proactive threat detection and risk prioritization. This allows investigators to focus on high-value targets and the most important information to the organization, thereby improving investigator efficiency.

Many organizations have accumulated various tools over time to safeguard critical information or assets, which makes it difficult to integrate different systems and technologies. With various back-end architectures and maintenance release dates, the ongoing administration of such an environment can be highly complex and expensive. Concerns about continuity of operations might deter software updates, leading to system performance issues or security vulnerabilities. IT professionals have enough challenges already. SAS offers a complete system based upon one common architecture, easing the deployment, maintenance, scalability and agility required to attain optimal security while keeping total cost of ownership low.

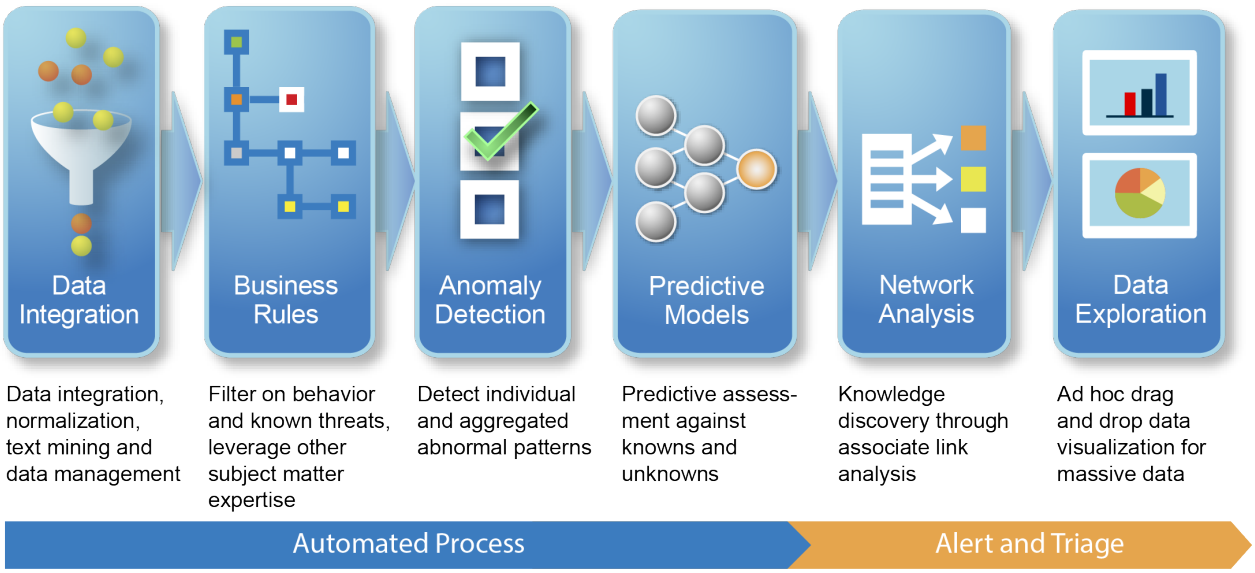

Figure 1: A hybrid analytical approach for insider threat deterrence.

The Foundation: Data Management

For many agencies, the process of integrating data (structured and unstructured), ensuring data quality (via cleansing and enrichment), and staging it in an analytics-ready environment is an arduous task. However, access to quality data is crucial for the output of sound and effective analytics. It's the cornerstone of every analytics endeavor with SAS.

Data Integration

Data integration is often underappreciated, but it can be critical to the success of analytical projects. Yet, manually bringing data together from multiple data sources with significant variations in format and context can take a great deal of time and effort, especially when sources contain unstructured data like text messages and emails – all of which can be useful for assessing risk.

The good news is that technologies now available can make it easier than ever to integrate unstructured data into analyses. For example, SAS facilitates effective integration of data and combines traditional entity extraction with point-and-click data management, sentiment analysis and other analytical methods. SAS can access and integrate virtually any type of data source, fuse data together, and enrich content so it's ready for analysis. Organizations can use SAS, for example, to access common data sources such as hiring information, physical security logs, SYSLOG (network) data, software compliance logs, and text logs (internal instant messenger or email communications).

Analytics Engine: The Hybrid Approach

SAS Analytics for insider threat deterrence is based on four distinct analytics domains that look at data from different perspectives. For example:

- Rules test all behaviors and activities against a predefined set of algorithms or business rules that can detect known types of risk behaviors based on specific patterns of activity or defined actions.
- Anomaly detection, such as clustering techniques, determines baseline behaviors for both individuals and groups and patterns of activity to define what's normal for each, measuring variations from the norm.
- Predictive modeling and data mining uses historical data or large amounts of transactions to predict future behavior and potential risks, as well as to detect new or emerging threat behaviors.
- Network or link analysis goes beyond data visualization to calculate the statistical significance between connections or transactions in the data and determines inferred relationships.

The following sections will explore the value of each analytical method in more detail.

Rules

Many security solutions rely on rules to generate alerts or leads – for example, when people cross a threshold by initiating a massive data download to an external device, or when someone makes multiple, denied database access attempts. The problem is that alerts are also generated for legitimate behavior, for example, when an employee forgets a password and keeps trying until access is denied. Security personnel can be quickly overwhelmed by an enormous number of alerts, which can drown out signals of real threats.

Applying multiple analytical methods simultaneously provides detection and identification of high-risk patterns of behavior and activity – even if they are novel or unknown – and presents them for triage. This hybrid approach significantly reduces the number of low-value policy deviation alerts (i.e., false positives) that can distract security professionals from the truly high-risk activities that warrant timely attention. This allows organizations to detect and address threats before losses occur, which better safeguards the organization and the individual from a life-changing and irreversible act.

Anomaly Detection

With anomaly detection, outlier behavior is identified through analyzing variables that identify behaviors and activity patterns worthy of investigative attention. Anomaly detection, such as clustering techniques, is critical to establishing the normal or typical activities of employees and their peer groups. Without an understanding of what is normal within an organization, it can be very difficult to determine what activities and behaviors are deviant and likely to pose risk.

> "The Hawthorne effect is named after what was one of the most famous experiments … in industrial history. It marked a sea change in thinking about work and productivity. … Hawthorne set the individual in a social context, establishing that the performance of employees is influenced by their surroundings and by the people that they are working with as much as by their own innate abilities."
> – The Economist[7]

7 The Hawthorne effect, *The Economist*, Nov. 3, 2008. economist.com/node/12510632

Using anomaly detection to identify outliers alone can present limitations. This approach by itself can generate too many false positives in instances where an employee's behavior changes significantly from known baseline behaviors, for example, if standard work times or intranet access patterns change. Used alone – outside of the hybrid analytics context – anomaly detection would incorrectly identify these changes as indicative of risk because it would fail to take into consideration additional context, for instance, that the employee was recently promoted to a new position for which these new patterns of behavior are appropriate.

Predictive Modeling

Predictive modeling uses historical data to create models of known threat behaviors and patterns. These models are then applied to new data to identify potential threats. Models can be run in real time or in batches, and they are especially useful in prioritizing risk across a low to high spectrum of potential risk probability. Those with a higher score are less likely to be false positives, which ultimately increases the rate of accuracy.

Predictive modeling has been very useful in detecting fraud. The richness of available threat-related data and the ability to incorporate known signatures of fraudsters into models have made these models more accurate and fraud investigations faster and more efficient. The detection of internal threats with predictive analysis is slightly different. For the known examples in historical data of prior bad acts, supervised models (based on the historical examples) can surface similar risk and even somewhat similar risk (referred to as semisupervised models) where detection is loosely based upon examples, but not completely without examples. In detecting internal risk where motives and tactics may not be definable, unsupervised models can be used without the known historical examples to look for and identify an unknown high-risk behavior. Similar to the behaviors models from the clustering algorithms, but different in that there is no "baseline" comparison, unsupervised predictive models find the "zero-day" vulnerabilities, ones that may not have been previously defined.

Network Analysis

Network analysis, also referred to as link analysis, is a key component of threat and risk detection and a cornerstone of the hybrid analytical approach. The value of network analysis for inside threat detection lies in its ability to map out relationships between people, activities, assets, communications or transactions that may not be evident at first glance. The massive amount of data that technology generates makes finding patterns or unusual connections challenging. "Soft links" (relationships that have statistical significance but no direct connection within the data) can imply relationships or patterns that are not likely to be identified in manual review or using less sophisticated methods.

Network analysis not only visualizes data in a relational manner, but also uses a risk scoring engine that triggers an alert generation process when risks surface in combination with the other analytical techniques. For instance, as shown in Figure 2, alerts are presented in a prioritized manner using a risk score (i.e., the aggregate of multiple analytical findings via the alert generation process). A user can click on an alert to drill down to the supporting details and expose findings that are significant to the risk and context to better understand the holistic behavior of a subject.

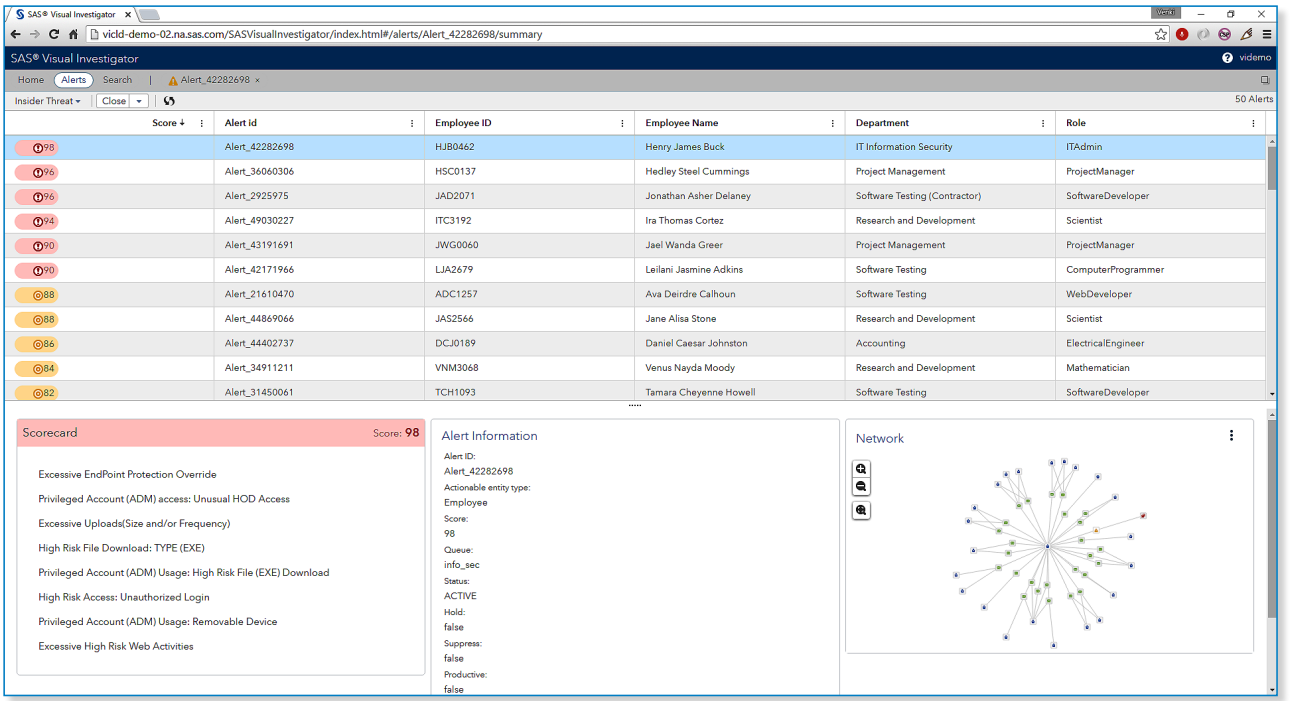

Figure 2: Prioritized risks using aggregated alert scoring and color-coded results.

The first alert shown in Figure 2 (the highest scoring alert) is regarding Henry Buck, a systems administrator who has been with the organization for about four years. The scorecard overview provides upfront data points to quickly provide the end user with a high-level understanding of the various attributes that compose the overall score. Additional views, such as the network diagram, can be customized for users or for the current investigative workflow.

Figure 3 illustrates the details of drilling down into an alert, and the in-depth evidence that surfaced the activity as a risk. The top left chart in Figure 3 shows a 30-day period of alert scores for the subject, Henry Buck, and his peer group, Systems Administrators. Review of the data in this visual format shows a correlation between the majority of Buck's activities and that of his peer group colleagues. This suggests to the investigator that the subject likely understands the organizational risk triggers, and that he may be conscious of his actions, careful to not attract unwanted attention. Further, some anomalies (on the 24th and 22nd) could imply his testing of the risk controls, and that other unusual behaviors over the past day have significantly deviated from the peer group and the organizational risk threshold.

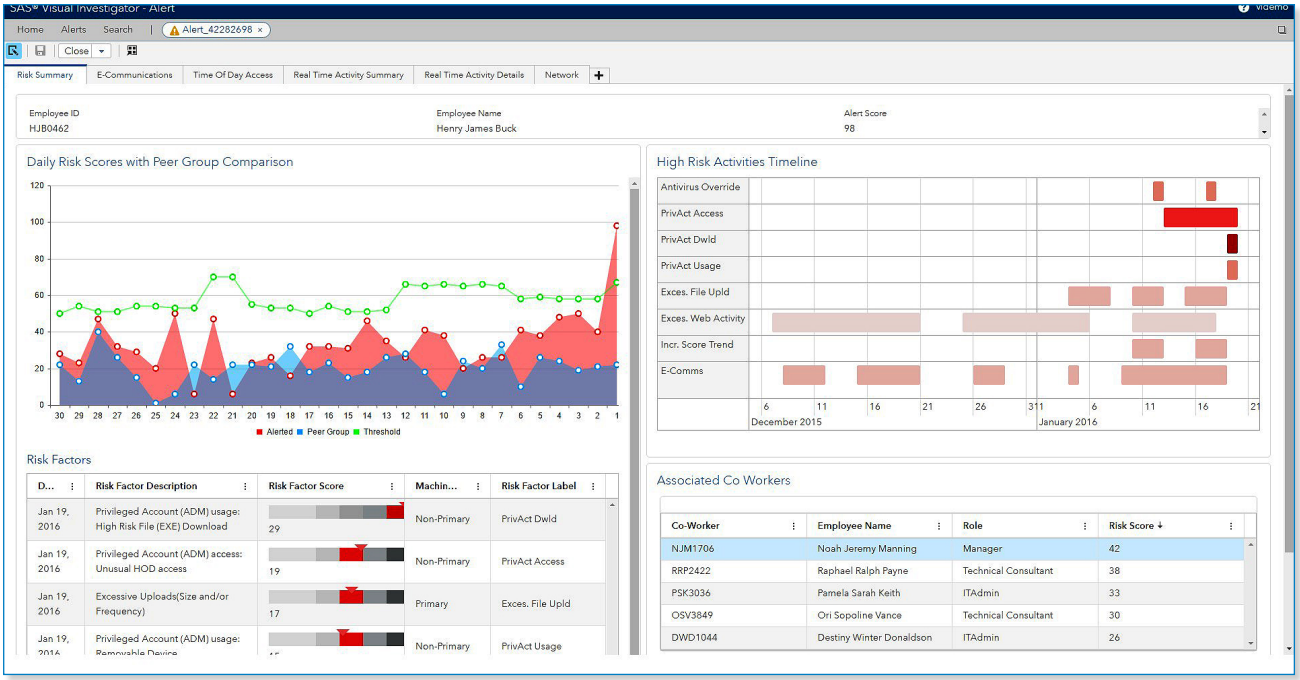

Figure 3: Visualizations of the context within the data describing the subject's high-risk behaviors.

Risk factors listed in the bottom area correlate to the specific daily score selected above, showing the data attributes that compose each daily score. This additional insight helps the user better understand what specific things the subject was doing; in this case, the subject was accessing privileged accounts (such as external file sharing) and using high-risk files (executables), had unusual time of day access, had excessive uploads, and used external or removable disks. Together, these activities paint a picture of activity that is unusual for his role or normal day-to-day behavior.

The right-side visualization in the interface shows the same data rendered differently in horizontal bar charts to show escalating behaviors of concern. This format helps the user see that over the last two months, the subject has consistently been engaged in high-risk web activities (defined by the organization as frequent visits to competitor websites and job searching websites), and has also frequently triggered high-risk alerts within communications analysis. Over the last three to four weeks, the subject has been uploading high-risk content. And recently, over the last two to three weeks, he increased unusual use of privileged accounts, overrode anti-virus software, and generated an increasing alert score. These details help investigators establish intent, identify gaps in data, and develop a case to determine if a formal investigation is warranted.

Having a systematic, risk-based, analytical approach provides a repeatable, auditable process that instills confidence in the results and facilitates communication with other stakeholders in the decision process. Should an investigation be warranted, the insightful content surfaced within SAS can be shared with the investigative team to provide a fast start in gathering missing elements, determining interviewing techniques, and for deeper forensic review.

Key Takeaways

The 2016 ACFE *Report to the Nations on Occupational Fraud and Abuse* analyzed 2,410 occupational fraud cases that caused a total loss of more than $6.3 billion.[8] Victim organizations that lacked anti-fraud controls suffered double the amount of median losses.

SAS' unique, hybrid approach to insider threat deterrence – which combines traditional detection methods and investigative methodologies with behavioral analysis – enables complete, continuous monitoring. As a result, government agencies and companies can take pre-emptive action before damaging incidents occur. Equally important, SAS solutions are powerful yet simple to use, reducing the need to hire a cadre of high-end data modelers and analytics specialists. Automation of data integration and analytics processing makes it easy to deploy into daily operations.

While hybrid analytics isn't a panacea for all insider threat problems (for example, it may not detect mental health/behavioral issues and external influences not evident within an organization), one thing is clear: Adopting hybrid-driven continuous evaluation and monitoring programs can significantly reduce insider threat risks. Why? Because used together, multiple analytical approaches are more effective than any single analytical approach. They act as a force multiplier, helping organizations do more with less – and do it faster and more efficiently – while improving overall security and reducing the risk of losses.

For hybrid analytics solutions to be successful, organizations also need well-established policies and executive support at the program level, as well as ongoing training initiatives. These efforts help organizations safeguard their most sensitive assets. They can act proactively and faster by reducing manual processes, ultimately mitigating losses rather than investigating them after the fact.

Learn More

For more information about SAS Security Intelligence, visit sas.com/vi.

8 *Report to the Nations on Occupational Fraud and Abuse: 2016 Global Fraud Study*. Association of Certified Fraud Examiners. s3-us-west-2.amazonaws.com/acfepublic/2016-report-to-the-nations.pdf

To contact your local SAS office, please visit: sas.com/offices