Red Flags Reimagined

Val LeTellier

DRAFT 05.25.22

A Tough Job Getting Tougher

The last few years have been particularly challenging for insider risk professionals. 'The trends are not our friends.' Employees are choosing permanent remote work, the 'great resignation' is overwhelming offboarding resources, nation states and criminal groups are getting bolder at recruiting employees to steal and ransom data, and COVID and political divisions are increasing employee stress, distraction, and disenfranchisement. To borrow from the cybersecurity 'CIA Triad' model, the *Confidentiality, Integrity and Availability* of our people, processes, and property are at risk.

And as reflected in the increasing number and costs of insider events, traditional countermeasures simply aren't up to the task. Specifically -- observable indicators are diminished by employees being 'out of sight, out of mind'. And network monitoring solutions only go so far, are complicated by remote work, and in general -- are reactive --- *not proactive*.

To illustrate our challenge, mentally put yourself in the chair of the insider risk analyst at a large organization. Each day begins fresh with the need to somehow identify a few potential bad actors from thousands of employees.

Paraphrasing former FBI Director Mueller, "There are only two types of organizations: those that have had a malicious insider and those that will have one. And even they are converging into one category: organizations that have had a malicious insider and will have another."

But it gets better – you also need to identify potential negligent or accidental insider risk. And you need to balance employee privacy, welfare, morale, organizational culture, and possibly even -- trusted workforce and zero trust strategies. And the consequences of a single malicious insider act can ruin your day and the organization. It's a high wire act. And none of these challenges are going away. It truly is 'evolve or die' time.

But enough admiring the problem. What can we do about it?

First things first: let's quickly address early warning, arguably the most critical element of insider risk mitigation, but often also the most neglected. Why? Because the problem is hard and complex.

To quote Marty Byrde from the television series Ozark, "As individuals, people are completely unpredictable. One person making one bet, I couldn't possibly tell you what they're going to do. But the law of large numbers tells me that a million people making a million bets --- that is completely predictable -- completely ordered." So, apply that to our challenge: insiders = individuals = unpredictability.

Capability, Motivation and Opportunity

We often speak of the insider's capability, motivation, and opportunity. Perhaps we should focus on our own – we are clearly motivated to stop insider attacks, but what about the capability and opportunity?

Opportunity

Fortunately for us, the insider's critical path toward action provides the opportunity to engage before an incident occurs. But generally, we don't maximize that opportunity.

Insider threat early warning programs often lack the necessary attention, expertise, funding, incentives, information-sharing, creativity, and programmatic approaches. Additionally, organizational cultures often undercut the effectiveness of early warning programs through denial, privacy concerns, lack of accountability and a cognitive bias toward technical cybersecurity.

And as you know, the move to remote work is particularly detrimental to the traditional red flag methodology. Because behavioral observation is a leading way that malicious insiders are discovered. And with many workers now only observed through the limited aperture of a computer screen, this countermeasure is largely lost.

Capability

In a perfect world, what would help? Quite simply -- efficient and effective early warning of at-risk employees. The automated identification of a very limited number of employees that require a closer look. The continuous evaluation necessary to remediate employee issues and create a trusted workforce -- so important now with so many people working remotely. But how to do we accomplish this?

Well, one way is by leveraging the advances in what I call our 'new operating environment'. Klaus Schwab of the World Economic Forum predicted that the Fourth Industrial Revolution would bring the "fusion of our physical, our digital, and our biological identities."

And this is happening. We see it every day in our lives and in the news. Data analytics connects dots that once took days or weeks to link - if it could be linked at all. And this fusion enables multiple surfaces to track, assess and even predict behavior in real time. The implication is significant for government and corporate security officials; new mechanisms and methodologies are now available to identify and mitigate risk.

And for insider risk professionals, we can use this algorithmic-fueled fusion to quickly highlight individuals and areas of concern. We can run behavioral, network, access, public data, and other feeds through link analysis and machine learning to create tailored advanced predictive analysis of thousands of employees in a few minutes. We can identify, sort, and stack malicious indicators into risk models that enable holistic continuous evaluation and zero-trust governance. Simply put, we can make insider risk mitigation smarter, more efficient, and more proactive.

But most importantly, we can provide 'decision advantage' to analysts and program managers. We can leverage internal and external data to create risk scores that highlight employees requiring analyst review. In the intelligence world, we call that tipping and cuing. But how do we create this decision advantage?

Let me propose five connected concepts. The first two concepts (balance and stakeholder buy-in) are meant to create the proper environment, and the last three (a holistic approach, the right data and advanced modeling) are meant to create the right process.

So, bear with me for a couple minutes – this is the meat of the sandwich.

Balance

First off, we need balance within our program environment. Balance between the risk mitigation we need, the employee welfare we seek and the employee privacy we must protect. As important as 'decision advantage' is to a program, it can't be at the expense of personal privacy. And privacy and trust are symbiotic. And make no mistake, the loss of employee trust is a sure-fire way to undermine a program.

So, while we've all surrendered varying degrees of our digital privacy to 'surveillance capitalism', we need to transparent in the program methods, processes, and goals. We need to show how we use anonymization, masking, generalization, and encryption to protect privacy. We need to evolve our marketing alongside our methodology and make insider risk mitigation less about threat reduction -- and more about employee welfare. We need insider risk mitigation to be viewed as an element of a positive security environment that: improves overall employee welfare and retention, maximizes business continuity, strengthens the organization's reputational value, and adds to the organization's bottom line.

That said, this is not a 'one size fits all' situation. Different organizational cultures merit different levels of privacy protection. This is particularly true between the 'cleared space' where employees hold security clearances and the rest of the employment market. To summarize, the insider risk equation is unique to each organization and has serious consequences if miscalculated.

Stakeholder Buy-in

And second, you need leadership and employee buy-in to the program. Not just to the *end goal* -- but also to the *necessary means* to accomplish that goal. So, take a moment and look at insider risk mitigation through the eyes of your organization's leaders and employees – how does it look? If this makes you uncomfortable, you have work to do.

The goal is to create a positive security culture. You can leverage your organizational culture and morale, and explain that the program is meant to provide early warning of employees who may need assistance from the organization. By doing so, the program will likely be viewed as positive rather than punitive.

Now that we've defined the necessary environment for our enhanced early warning program, let's get into the details. We've already established the critical importance of intercepting at-risk employees in the earliest stages of the critical path.

A Holistic Approach

As we've learned, the same 'root causes' -- personality predisposition and critical events --- can result in self-healing or self-harming responses by different people. And for those seeking to harm, there are different forms they may choose as action (theft, sabotage, violence, suicide, etc.).

So, he third concept is a holistic approach. One that takes into consideration the individual *and* their mental, emotional, financial, physical, virtual, chronological state and environment. Specifically, we need a "whole person" and 'whole threat' approach.

To me, 'whole person' is contextual and psychosocial, using personality, environment, and precipitating events to identify insider risk. And 'whole threat' addresses the common root causes that result in different forms of attacks (data theft, fraud, sabotage, violence) – and in all domains (cyber, human, and physical), narrowing the attention to critical materials, data, and processes and those with access to those items.

Combined, the whole person and threat approach focuses an organization's limited resources on its: most sensitive holdings, the insider personalities meriting greatest concern, the precipitating events that can turn those personalities into harmful actors, and the corresponding indicators that highlight the need for closer inspection.

The Right Data

But to pursue a whole person and whole threat approach, you need the right data. To quote former Hewlett Packard CEO Carly Fiorina, "The goal is to turn data into information and information into insight'. But first you need the data. And obviously, the better the data, the better the analysis and the more accurate the risk scoring.

To get the best data, we need refined research and understanding of which indicators are statistically proven against the progression of different insider types along the critical path. We need behavioral psychologists, insider risk analysts and data scientists to help us find the right combinations of data capable of highlighting the disparate indicators taken from thousands of cases. We can't have 'garbage in, garbage out' --- we need 'quality in, quality out'. We need a data ecosystem relevant to the insider risk equation: human behavior, network activity, physical accesses, and the context to make sense of it.

As examples of relevant internal data sets, think of employee:

- Access to sensitive materials.
- Network behavior.
- Facility access.
- Sensitive materials exfiltration.
- HR issues (performance plans, resignations, terminations, etc.)

For external sets, think of an employee's public data records, particularly those reflecting:

- Financial and legal life stressors.
- Extremism.
- Foreign connections.
- Conflicts of interest, etc.

Advanced Risk Modeling

And finally -- we need advanced risk modeling. This is where the 'magic' happens, this is where a 'digital twin' is created, and This is where we get into the head of the insider and understand what sets them off, how they would plan and act. This is where advanced analytics and fusion technologies eliminate the spaces between data points and connect dots that once took days or weeks to link — if they could be connected at all.

This fusion is enabled by a tailored suite of algorithms and machine-learned analysis that churns through internal and public records and live sensor feeds. And continuously develops employee risk scores based upon

sensitive material access, personnel and security issues, physical security, life stressors, network behaviors, etc. This allows 'risk triage' of large employee populations, and a manageable number of cases for analyst attention. But to do this efficiently and effectively, we need to understand the insider profiles most relevant to our organizations. And we need to develop and automate a watchlist of the most relevant trip wires.

The Good, the Bad, and the Ugly

But before we go all in, let's take a clear-eyed look at the potential consequences of applying holistic continuous evaluation to an organization:

We've covered a good bit of the good:

Preempting costly incidents. But also fostering a positive security environment, improving employee welfare and retention, maximizing business continuity, strengthening the organization's reputational value, creating risk scores for zero trust architectures, and quantifying trusted workforce metrics.

But there is also some bad:

Leadership approval. Employee monitoring of any type is often categorized by leadership and employees as a form of 'big brother' surveillance that shouldn't exist in their organizational culture. Covid and the 'Great Resignation' complicates matters; after losing a direct view of their employees, many organizations moved toward digital surveillance to ensure productivity. But as firms implemented this, many of them experienced increased employee turnover. And with a global shortage for workers, the tolerance for anything that diminishes employee retention is now understandably very low. So, promoting this approach may be an uphill battle. That said, leadership often quickly abandons this view once they experience a 'significant emotional event' with a costly event, it's a fact of life that some organizations are simply not prepared to do the needful.

And finally, there is the ugly:

Shadow IT. While some employees respond to digital monitoring by finding a new job, others skirt official corporate networks, tools, and safeguards in favor of their own devices and infrastructure, which in turn can create insider events. The increased use of shadow IT is truly concerning, particularly as its often accompanied by cutting corners and neglecting policies and procedures – all so the employee can save themselves time and effort. A good example is forwarding sensitive data to a personal device or account so it can more easily be edited or printed. By doing so, that information then rests unprotected on a personal device, outside the organization's (and maybe employee's) control and is highly vulnerable for theft or exploitation.

Summation

Remote work creates new attack vectors and raises monitoring challenges, Covid increases employee stress and more negligence and 'the Great Resignation" increases the statically proven high data theft rates accompanying employee departures. Diminished employee visibility, increased employee stress, and overburdened monitoring systems create a conducive environment for harmful insider action.

Meanwhile, most existing solutions only take us so far. They are cyber-focused: log centric, singularly focused on network anomalies and retroactive, and are failing us. Additionally, nation states and organized criminal groups are employing sophisticated espionage and social engineering tradecraft to steal credentials.

It's been high-stakes chess game, even before the recent rise of ransomware and state-sponsored cyber warfare, and we need to be thinking several steps ahead of the challenge. Never have we more needed strong and modern insider risk tradecraft.

And to enable that modern tradecraft, we need to harness modern technology to create proactive continuous evaluation that enables early engagement of at-risk employees, remediation of toxic situations, and preemption of costly and life-threatening incidents while maintaining a focus on employee welfare and employing end-to-end privacy.

We need to be constantly objective, start with the premise of innocence, and use observation as a starting point for further information collection (versus judgement). And finally -- we need to predetermine status that if achieved, will trigger an end to an investigation (reverse tripwires). This will enhance morale, boost stakeholder buy-in and help create a positive security culture.

Theory, Meet Application

To quickly illustrate how this would work, please see (include link) as to how this methodology could be applied against classic insider types.

Val LeTellier ran security, intelligence, and counterintelligence operations as a State Department Diplomatic Security Special Agent and CIA operations officer. Twenty years penetrating foreign intelligence targets and recruiting sources provided him an intimate understanding of the psychology of insiders. Following government service, he co-founded a cyber security firm that combined CIA HUMINT and NSA technical expertise for insider risk vulnerability assessment and countermeasure design. He now helps develop innovative security and intelligence solutions for the new operating environment created by the Fourth Industrial Revolution. He presents nationally on insider risk, and holds an MS, MBA, CISSP, CEH, PMP, RTT and ITVA.