
How machine learning is catching up with the insider threat

Received (in revised form): 26th May, 2017



Jamie Graves

is a data security and enterprise software entrepreneur and is the founder and CEO of ZoneFox. He attended the prestigious Ignite course at Cambridge University's Judge Business School, and the Entrepreneur Development Programme at MIT. He has a PhD in Computer Science, extensive security and digital forensics experience, and was recently recognised as the 'Champion of Champions' at the inaugural Scottish Cyber Security awards for his contribution to the industry. ZoneFox protects business-critical assets, data and IP and has a proven track record in protecting reputation, sales revenue and competitive advantage for its growing client base.

CEO, ZoneFox, 3 Lady Lawson Street, Edinburgh EH3 9DR, UK
Tel: +44 (0)845 388 4999; E-mail: j.graves@zonefox.com

Abstract The insider threat poses a unique cyber security challenge. When it comes to meeting this challenge, the type of 'standard' threat detection toolsets currently deployed by organisations tend to be inadequate. This paper aims to show how and why machine learning capabilities can help organisations to reduce these inadequacies, providing an essential extra element of protection. The paper explores the reality of the insider threat, illustrating that while the possibility of a malicious actor cannot be discounted, this threat is much more likely to arise through carelessness, inadvertence or lack of understanding. A focus on best practice and clear policies must always be part of the solution — backed up by threat detection tools. The paper explains the problems that can arise with such tools, including the delays and inaccuracies that can arise with configuration and updates. With its focus on *behaviour* (as opposed to reliance on *signatures*), it examines how machine learning is able to determine 'usual' activities and flag up events that fall outside of the 'usual', and looks at the benefits this can bring to cyber security teams, in terms of ability to detect as wide a range of abnormal activities as possible, improved visibility, more accurate insights and better use of resources.

KEYWORDS: machine learning, insider threat, user behaviour, UEBA, forensics analysis

Insider threats abound nowadays, always changing, many hiding in plain sight. The reason that the insider threat gives us so much concern is that it doesn't always take a malicious actor to perpetrate acts associated with the insider threat.¹ Sometimes the insider threat stems from a configuration error, sometimes a lack of user education; inappropriate permissions often contribute to the insider threat.

While the best way to circumvent the insider threat is to establish policies, standards

and best practices to help secure your environment and provide optimal controls, the process to do so is very arduous. Luckily, technology now exists that can help fill in some of the gaps while we get our ducks in a row. Detection technologies that leverage machine learning can act as a buffer to help find threats to which we were previously oblivious, helping us catch up with the insider threat. Let's take a look at how organisations are affected by the insider threat, and how machine learning is helping us catch up.

THE INSIDER THREAT — AND WHY IT'S A THREAT

When some think of the insider threat, they think of a sleeper cell working within their walls, just waiting until the time is ripe to steal all of their organisation's data or plant a backdoor through which they will be able to access sensitive data whenever they'd like. In reality, however, the insider threat is generally much more mundane. Inadequate role-based access control, user education or basic cyber security controls can result in higher levels of insider threat. Add to this the fact that most cyber security controls are network-based, then much of the activity a user might take to exfiltrate data — unwillingly or otherwise — may go undetected.

For example: Carl, who is rather careless, decides that he is going to work at home tonight to show the boss that he's willing to go the extra mile to get stuff done. Carl loves his brand new, superfast Mac Pro that he uses in his home office, so he decides to put the specification document (which happens to be intellectual property) on which he is working onto a USB stick so that he can take it home with him and work after his kids go to bed. On his way home after a drink with the team, Carl happens to drop his USB stick on the ground in the pub car park. Another car park user notices and picks up the stick. Now Carl's company IP is in the hands of a stranger who didn't even have to hack into its network for what may be a set of extremely valuable data, ready to go to the highest bidder. While it may seem a bit far-fung, you can see how elements of this scenario are not so much — several high-profile cases of missing or lost USB sticks over the years can attest to this.² But what if you could detect the use of Carl's USB drive, and the copying of confidential files to said drive? Well, there have been tools developed to do just that.

CURRENT TOOLSETS FOR HANDLING THE INSIDER THREAT

Many tools do exist, currently, that can help detect and mitigate the insider threat, the

main ones being SIEM (Security Information and Event Management); DLP (Data Loss Prevention) solutions and IPS (Intrusion Prevention System).³

These tools may not be labelled as insider threat detection tools, but in essence they do just that. Data loss protection solutions provide a layer of detection and prevention for data exfiltration or destruction. Intrusion prevention systems help us detect malware or other attacks that may compromise our endpoints, providing malicious actors access to our information assets. Web content filters help ensure that users do not try to upload data to cloud storage services or send it as webmail attachments. Endpoint protection solutions are now generally kitted out with firewall or host intrusion prevention system functionality, as well, to help layer up the defences. Throw a SIEM into the mix, and now you've got a central location for all of your logs, as well as the ability to aggregate and correlate the events generated by each sensor and provide alerting based on the results of said aggregation and correlation. Looking at this toolset, things seem to be fairly locked down, right? Well, not entirely.

One big problem with the current toolset is that it relies heavily on human interaction being accurate. Humans need to develop the configuration for the tools, and that configuration generally stays pretty static throughout the lifespan of the tech. If new attacks arise, then tool manufacturers must release new signature updates to detect or prevent these attacks. In most cases, a new signature, plug-in etc. must be tested in a lab environment to ensure that it does not adversely affect your users, then pushed to production before it can actually do its job. The 24–48 hours (or more) required to test these signatures requires human interaction. Other technologies require that you have analysts, engineers or other folks who work on the tools creating policies, rules or filters to detect malicious activities. Configuring static rules to detect credit card numbers are great, but DLP solutions that live or

die by their policies simply cannot catch everything unless you have a watertight data classification scheme.⁴

Another significant issue with today's approach is that many of the tools used to mitigate the insider threat can cause users some inconvenience. SIEM tools alone can be complex and resource-intensive to both install and run, and DLP — often as complex as SIEM — comes with all kinds of baggage, from user disruption and privacy issues to the aforementioned requiring visibility of unstructured data. If any shortcomings in technological solutions affect productivity or user satisfaction, you can bet your boots that the technology will be removed from production, or the offending features disabled.

Again, in order to ensure that any features do not adversely affect production environments, you need a team that can test the tools and features. This can take a long time to get just right, and in many cases the technology simply doesn't get used to its fullest potential. When we also factor in tight deadlines, human error and budget, the effectiveness of these tools can be greatly reduced in the name of saving some coin.

HOW CAN MACHINE LEARNING HELP US CATCH UP WITH THE INSIDER THREAT?

In order to better detect and contain the insider threat, we need to first understand that it's dynamic in nature. You can wrap a high-level definition around it, but you cannot use a general rule or security policy to encompass it. We need to be learning on the fly, understanding what our users are up to on a day-to-day basis. Keeping track of that type of activity requires data. Lots of data. While we can employ tools such as those mentioned above — SIEM, DLP, IPS — they provide us logs as data, logs that need to be normalised, identified and filtered down to the nuts and bolts that might be useful to security analysts. What's more, said

nuts and bolts should be delivered in real time, to be actioned and ensure damage is mitigated.

SO WHAT EXACTLY IS MACHINE LEARNING?

Machine learning is essentially a cross between data analysis and automation. Solutions that leverage machine learning, ranging from user and entity behaviour analytics to self-driving cars, are built on algorithms that focus on constantly ingesting data, continuously analysing it, and then making decisions.⁵ These algorithms create a sense of thinking, or learning. Whether it's identifying a user in your network who has performed activities that do not fit in with the profile that the solution is created, or identifying road construction and finding an appropriate detour, machine learning is constantly ingesting data and producing actions to provide a fast, effective way forward.

Machine learning has been around for a while, but has only recently been applied to big data to create advanced cyber security analytic capabilities. The benefits — and potential benefits — are myriad. By leveraging data, not just logs, to better understand user behaviour (also known as user and entity behaviour analytics, or UEBA), we can determine our users' 'usual' activities and flag any behaviour that falls outside of the 'usual'. Applying statistical data models and metrics can make a world of difference when detecting insider threats, for a few simple reasons.⁶

MACHINE LEARNING DETECTS MORE THAN JUST THE BAD STUFF

The main goal of machine learning is to understand what your users are doing; for example, monitoring which file shares or other assets are being accessed regularly, which applications are being run on a regular basis, and which user logs on to a system can

tell you a lot about your users' behaviour. For instance, Sandra has a basic routine that she follows every day: she accesses the accounts payable report every morning on the finance file share; receives e-mails containing various invoices, requests for credit and money transfers throughout the day; and saves the accounts payable report to the same spot with a different filename every day. If one day she attempts to access files on the engineering share, this falls outside of her usual activity. Maybe it's nothing, but maybe it's something. Wouldn't you like to know either way? Well, the baseline your machine learning-enabled UEBA tool will be able to detect and report on such activity. If you have a developer who appears to be doing way more downloading of code than uploading, UEBA can detect this behaviour and flag it as suspicious.

SPEED IS ESSENTIAL, MANUAL CONFIGURATION IS NOT

With machine learning, your engineering team does not have to spend hours on end analysing traffic patterns and creating rules and filters to provide precise alerts based on an event that must meet several criteria to be classified as an attack. One of the huge benefits provided by machine learning is that the platform itself will do the analysis and create the rules, leveraging the aforementioned statistical models and probability metrics. Simply put, machine learning platforms take historical and real-time user data to help form their own opinion, rather than be indoctrinated by various rules — most likely dictated by a cyber security standard or best practice, rather than by past events occurring on the network. Within hours, potentially, of placing sensors in your environment, your UEBA platform can form its own opinions and start detecting any anomalies on the network. Another great feature of most UEBA platforms is that they can be deployed relatively easily.⁷ You can generally skip

the months of planning, as most of these platforms sit off a tap or span port and just need to see copies of the traffic; they do not need to sit inline and potentially disrupt the traffic flow.

REACTIVITY IS FUTILE; PROACTIVITY IS KEY

Currently, we live in a primarily reactive state when it comes to dealing with the insider threat. Machine learning provides a level of proactivity; what with predictive analytics and probability metrics, we can start to see a little bit into the future — one of the infinite possible outcomes, anyway. How proactive your cyber security capability is can make the difference in stopping the insider threat before it gets a chance to pilfer your data out, let the bad guys in, or do any other lasting damage to your data, assets or reputation. In a reactive world, we are constantly waiting for the next incident to occur, while hoping that we can actually see it. With a machine learning solution that provides a new level of proactive security monitoring, we can identify indicators and provide warnings to security analysts so that they might be able to preempt a malicious attacker or unwitting user and stop the incident from happening in the first place.

MACHINE LEARNING LEAVES EMOTION OUT OF THE EQUATION

The lack of emotion in machine learning UEBA platforms is somewhat double-edged. Objectivity is great when it comes to detecting and responding to insider threats. While the maths behind machine learning can be pretty complicated, it's at the very least non-partisan. UEBA platforms do not pick favourites. They do not have friends, and do not believe that they can let some folks off with violating cyber security policy. Machine learning does not slow down due to distraction, anxiety or a bad break-up. UEBA platforms will give you the same results based

on the same data set — and that's great. The downside to such objectivity is that the platform does not understand whether or not a given user is having a bad day, simply mis-clicked, or was just curious and bears no ill will. So, before you fire your security analysts and use their salaries to procure a UEBA solution with machine learning, keep in mind that you will need their humanity to help the UEBA platform discern whether the alerts actually signify a malicious insider, an unwitting insider or a configuration error. While these may all amount to the same type of alert within a UEBA solution, the fact is that each of those incidents would need to be treated differently.

Machine learning still requires human interaction, just less than today's mainstream monitoring solutions. Although one of the greatest benefits of machine learning solutions is their independence from human interaction, they cannot reach their full potential without some form of help.

A UEBA solution that uses machine learning doesn't require you to create policies, filters or rules in order to detect potentially malicious behaviour; it uses its own statistical models to deduce whether or not a user is doing something out of the ordinary. Human interaction is required, however, to tell the UEBA solution whether or not it's correct. Adding a thumbs-up/thumbs-down (to put it simply) result to the equation tells the UEBA whether it's on the right track, or if it's throwing up false positives and needs to recognise the action that is under scrutiny as legitimate.

The UEBA solution with machine learning provides more accurate events coming into your cyber security operations centre.⁸ The net result of this added accuracy is far less time spent by your security analysts actually analysing events (and potential events) to be added as criteria to a SIEM. The difference here is prediction versus reality; in this situation, what would you prefer?

- UEBA with machine learning provides more accurate insights than conventional security solutions.⁹
- Human input into UEBA solutions with machine learning can help the solution provide even more accurate results the next time around.
- Machine learning leverages statistical models and probability to provide alerts, instead of analysts manually creating scenarios which will trigger alerts.
- Solutions that use machine learning can still provide false positives — analysis is still required!
- Machine learning solutions can save hours of analyst time creating rules and filters.

Proactivity is key when it comes to cyber security, but we still must maintain the ability to react. Machine learning is great because it's proactive, but while proactive detection provides a great advantage, it's nothing without the ability to react appropriately to any security events or incidents that may arise.

In addition to machine learning, we also need common sense. While common sense is not all that common, we can help build it into our cyber security practice through the introduction of processes and standards. An incident response plan — especially one that specifically takes into account the insider threat — is one of the keys to success in cyber security.¹⁰ If your machine learning solution picks up a potential insider threat, how will your response team react? Will they light their torches and grab their pitchforks, or will they have a calculated, uniform response that includes HR, management and IT? Proactivity gives us the jump on any potential attackers, whether inside or out. The inability to properly respond, however, will still leave something to be desired. In short, ensure that alongside the next-generation machine learning solution you implement within your environment resides a tried, tested and true incident response plan.

BUT THE MACHINES CANNOT TAKE ON ALL THE LOAD!

Machine learning provides us with unbiased, objective results, which can be truly invaluable in an investigation. That said, the forensic data provided by a machine learning solution does not replace good old-fashioned forensic analysis. Although the output provided by UEBA solutions is generally accurate, gathering supporting evidence is a must. Collecting system, network or application logs is necessary, and let's not forget the nitty gritty of forensic analysis. You will still need to gather forensic images of suspect system hard drives and memory dumps in order to determine whether a user has performed the activities in question, such as deleting data or copying it to a USB drive.

The benefit provided by a UEBA platform that leverages machine learning is that you may not have to sift through all of the collected data. Your UEBA solution should provide a giant glowing arrow pointing directly to abnormal activities; all you should be required to do as a forensic analyst is to follow the arrow and dig deep to find the data in question. This could save days, even weeks, of analysis up-front.

CAN FORENSIC ANALYSIS BE AUGMENTED WITH MACHINE LEARNING?

While security monitoring solutions that leverage machine learning can find events which may escape conventional detection measures, can it help augment digital forensics? There is at least one use case for machine learning, for sure. If you are performing multiple analyses on varying sets of data pulled from several endpoints, building up a database of findings would be beneficial. Leveraging big data and machine learning enables forensic analysts to store vast amounts of objective data about their findings into a repository and detect abnormalities in an automated fashion. For example, if you have a user who has been

up to no good, or has been compromised by a malicious actor, you can take memory and filesystem dumps from your forensic image, normalise or index the data, and have your machine learning solution digest the data — the end result being a potential for faster, objective, more accurate forensic analysis. User validation comes in the form of searching for the flagged data, verifying and providing input back into the system to help identify false positives. The same model for security monitoring could be used, just in a historic fashion instead of real-time.

CONCLUSION

The insider threat is finally being recognised as one of the greatest threats to an organisation's data, intellectual property and assets. Not that it was ignored previously, but there was a sense of helplessness, or false bravado around handling the insider threat. Ultimately, we had been unsuccessful in providing a solution to help response efforts around the insider threat — that is, until UEBA platforms with machine learning capabilities came to the rescue. While these platforms have not yet reached ubiquity in the modern organisation, much traction has been gained in a short time, and the future looks bright. With many cyber security vendors, particularly those of SIEM and data analytics platforms, getting on the machine learning bandwagon, it should soon be part of any cyber security operation's arsenal for combating the insider threat. Although the human element is still required to analyse events, confirm that they are malicious and provide further forensic detail, machine learning has the potential to take us leaps and bounds ahead when it comes to detecting, analysing and responding to the insider threat.

References

1. Sculze, H. (2016), 'Insider Threat Spotlight Report 2016', Crowd Research Partners, available at <http://crowdresearchpartners.com/wp-content/>

- uploads/2016/09/Insider-Threat-Report-2016.pdf (accessed 30th May, 2017).
2. 'At least 1,000 government laptops and flash drives reported missing since 2015', *Guardian*, 21st December, 2016, available at <https://www.theguardian.com/politics/2016/dec/21/at-least-1000-government-laptops-and-flash-drives-reported-missing-since-2015> (accessed 30th May, 2017).
 3. Litan, A. and Carpenter, P. (December 2016), 'Best Practices for Managing Insider Security Threats', Gartner, available at <https://www.gartner.com/doc/3418831/best-practices-managing-insider-security> (accessed 30th May, 2017).
 4. Wikipedia, 'Data Loss Prevention Software', available at https://en.wikipedia.org/wiki/Data_loss_prevention_software (accessed 30th May, 2017).
 5. Wikipedia, 'Machine Learning', available at https://en.wikipedia.org/wiki/Machine_learning (accessed 30th May, 2017).
 6. Bussa, T., Litan, A. and Phillips, T. (2016), 'Market Guide for User and Entity Behavior Analytics', Gartner, Gartner, available at <https://www.gartner.com/doc/3538217/market-guide-user-entity-behavior> (accessed 30th May, 2017).
 7. Kim, E., Kish, D., Litan, A., Contu, R., Carpenter, P., Deshpande, S., Pingree, L., Ahlm, E., Heng, J. and Gardener, D. (February 2017), 'Market Insight: Security Market Transformation Disrupted by the Emergence of Smart, Pervasive and Efficient Security', Gartner, available at <https://www.gartner.com/doc/3592617/market-insight-security-market-transformation> (accessed 30th May, 2017).
 8. Arsene, L. (September 2017), 'How Machine Learning For Behavior Analytics & Anomaly Detection Speeds Mitigation', Dark Reading, available at <http://www.darkreading.com/partner-perspectives/bitdefender/how-machine-learning-for-behavior-analytics-and-anomaly-detection-speeds-mitigation/a/d-id/1327830> (accessed 30th May, 2017).
 9. Kim, E., Kish, D., Litan, A., Contu, R., Carpenter, P., Deshpande, S., Pingree, L., Ahlm, E., Heng, J. and Gardener, D. (February 2017), 'Market Insight: Security Market Transformation Disrupted by the Emergence of Smart, Pervasive and Efficient Security', Gartner, available at <https://www.gartner.com/doc/3592617/market-insight-security-market-transformation> (accessed 30th May, 2017).
 10. Bromiley, M. (June 2016), 'Incident Response Capabilities in 2016: The 2016 SANS Incident Response Survey', SANS Institute, available at <https://www.sans.org/reading-room/whitepapers/incident/incident-response-capabilities-2016-2016-incident-response-survey-37047> (accessed 30th May, 2017).