2018 TAG CYBER SECURITY ANNUAL

VOLUME 2

INTERVIEWS WITH CYBER LUMINARIES

Expert Advisory Research

Dr. Edward G. Amoroso
Chief Executive Officer, TAG Cyber

September 2018

Design – TAG Cyber LLC
Finance – M&T Bank
Promotion – TAG Cyber LLC
Facilities – WeWork, Fulton Street Station, New York City
Administration – navitend
Research – Liam Baglivo, Matt Amoroso, Miles McDonald
Lead Author – Dr. Edward G. Amoroso

TAG Cyber LLC
P.O. Box 260, Sparta, New Jersey 07871

Security experts and practitioners must recognize that best practices, technologies, and information about the cyber security industry and its participants will always be changing. Such experts and practitioners must therefore rely on their experience, expertise, and knowledge with respect to interpretation and application of the opinions, information, advice, and recommendations contained and described herein.

Neither the author of this document nor TAG Cyber LLC assume any liability for any injury and/or damage to persons or organizations as a matter of products liability, negligence or otherwise, or from any use or operation of any products, vendors, methods, instructions, recommendations, or ideas contained in any aspect of the 2018 TAG Cyber Security Annual volumes. The opinions, information, advice, and recommendations expressed in this publication are not representations of fact, and are subject to change without notice. TAG Cyber LLC reserves the right to change its policies or explanations of its policies at any time without notice.

September 7, 2017

To the Reader:

Interviewing the cyber security luminaries included in this *2018 TAG Cyber Security Annual* was a thrill for me on par with what a political scientist might experience interviewing world leaders. My hope is that the pure joy of learning afforded by these capable and successful cyber security experts comes through in the narrative and transcripts. As any interviewer will attest, the goal is for you the reader to feel like you were seated *right there* during the interview, learning from the insights and views of these fine security experts.

As is always the case with our work at TAG Cyber, we humbly reached out directly to these experts hoping they would be willing to share. We are not a massive company with analysts in every corner of the globe; rather, we are a small start-up working round the clock *trying harder than the other guys* to bring the best cyber security analysis to enterprise teams. This volume of interviews is hopefully evidence of our sincere pursuit. We stopped counting the number of hours that went into its production.

Editing down the interviews this year was harder than in our previous year, perhaps for no reason other than our subjects seemed more comfortable with what we are doing at TAG Cyber. Last year, I noticed slight reservation at times that we might quote-out-of-context or highlight-controversial-stuff in our interviews. Now that our focus on sincere learning is more obvious, and we've built up some reputational trust, our interview subjects were more relaxed. Everyone seemed to talk more this year.

Our advice on using this volume is simple: You can read it from start to finish, but recognize that the order is reverse-alphabetical (nice going, ZeroFox) for no reason other than convenience and fairness to last year's non-reverse-alphabetical approach[1] (nice going, Agari). Maybe next year, we'll do a random scatter. As an alternative to reading this volume from start to finish, perhaps you might use it as a reference guide on your virtual e-shelf to augment your understanding of a given area or vendor.

Regardless of how you use the volume, we are honored that you are spending time with our materials. Every word of every sentence of every page was written with one goal: To be of sincere assistance to the women and men who protect our world's systems and infrastructure from cyber attack. These folks are the unsung heroes of the technology revolution, and without them, our lives would be a pile of chaos. If you have any of these cyber defenders working for you, then please give them a raise.

I hope you all enjoy and learn from this volume.

Dr. Edward G. Amoroso
Chief Executive Officer, TAG Cyber LLC
Fulton Street Station on Broadway

[1] By the way – reverse alphabetical ordering is harder than you think. Try counting backward from Z right now in your head, and you'll see what I mean.

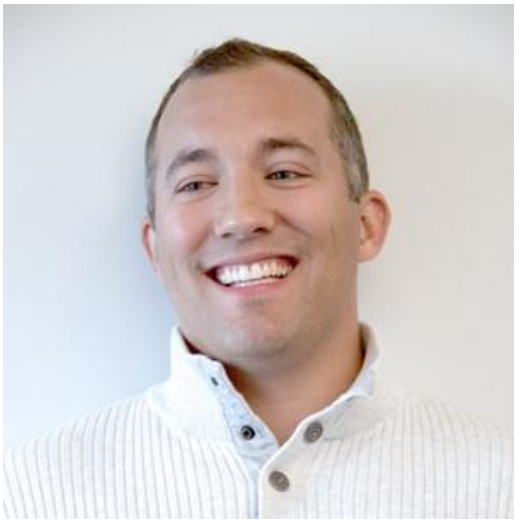

Securing Social Media Usage for the Enterprise

Threats to social, mobile, and digital platforms represent a significant new attack vector for the modern enterprise.

James C. Foster, CEO of ZeroFOX

It should come as no surprise that social media, mobile, and cloud platforms such as LinkedIn and Facebook represent new attack vectors for the enterprise. Most companies *rely heavily* on these social platforms in their marketing, messaging, and outreach programs for employees and customers. As such, new cyber security issues for these services have emerged and require immediate mitigation. Without such protection, corporate brands and reputations can be degraded. James C. Foster, CEO of ZeroFOX, knows quite a bit about this new area of enterprise security and he sat down with us to share his views on current and future trends.

EA: What are the types of risks that your team is observing on social media platforms?
JF: Anyone who has been watching the news lately understands the potential risks that exist on social media. These are incredibly powerful platforms that create communities and help people and businesses connect and share, but because social media is mostly unmoderated, the possibility emerges that misrepresentations can be made. This can have minor consequences when the activity is isolated, but it can have major consequences if it involves a business or many people. Reputations for organizations are no longer just maintained based on business activity, they are also now affected and influenced by external social, mobile, and other digital platforms. This is what drives risk.

EA: Do mobile platforms and app stores introduce risk to the enterprise?
JF: Anywhere spoofed accounts can be set up outside the perimeter, digital risks emerge. Since app stores are included in this category, they should be monitored to ensure that new risks to an organization have not been created. But it's not just app stores – it's any social, digital, or collaboration platform, and this includes Pastebin, Facebook, Reddit, and other popular forums. These all require monitoring from a digital risk perspective.

EA: Are there mitigations that can be performed when an enterprise is experiencing digital risk?

JF: Organizations are digital now and risks that target an organization on social, mobile, or collaboration platforms can be mitigated in many ways. Luckily, enterprise teams can mitigate the effects of phishing campaigns, customer scams, fraudulent accounts, and many other threats on social and digital platforms. We support this process through accurate monitoring, which is the first step to identifying any security problem, combined with advanced automation, which helps us ingest, analyze, and remediate malicious risks for our customers. We work closely with social media and other digital platforms to take down risks, content, and profiles that violate terms of service.

EA: How does your team's platform work? Are there experts working behind the scenes to assist in the risk monitoring function?
JF: The ZeroFOX Alpha Team is the only research team in the world that is dedicated to the identification of emerging threats and risks on social media and digital platforms. Our researchers are active in the security community, helping to bring down large-scale campaigns that affect everyone. For example, we might analyze tens of thousands of impersonator profiles to identify trends. Similarly, we might lead the investigation into Spam botnets spreading fake news, porn, and other unwanted content.

EA: What are your predictions regarding social risks to the enterprise in the coming years?
JF: We have always seen that the bad actors go where the vulnerabilities are – and always target unprotected people, businesses, and data. Since people, businesses of all sizes, and our most up-to-date data are on social media and digital platforms in huge numbers and growing, these challenges are not going away any time soon. Like email before it, social media is the number one form of communication and attackers will continue to be motivated to target all of us where we communicate to get to us and our data.

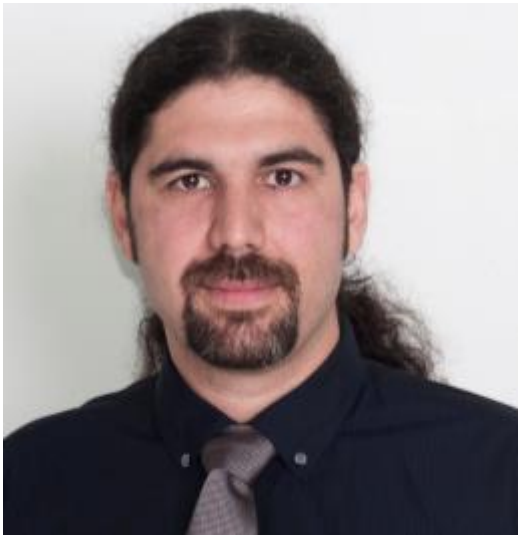

Unidirectional Gateways for ICS Cyber Security

Ensuring one-way dataflow between ICS devices and support infrastructure offers powerful separation control.

Lior Frenkel, CEO of Waterfall Security

Industrial control systems are just as susceptible to cyber attacks as other aspects of modern technology infrastructure. The problem is that the consequences of attacks on the operational technology (OT) associated with ICS can be more intense than the types of issues that can result from traditional IT risk. Safety and life-critical implications often arise, for example, in OT/ICS security scenarios. Lior Frenkel, CEO of Waterfall Security sat down with us recently to share the basis for his company's Unidirectional Gateways and associated technologies as powerful means for optimizing protection solutions for industrial control.

EA: What is meant by a unidirectional gateway?
LF: A Unidirectional Security Gateway is a technology that adds a physical layer of cyber security to the industrial network perimeter to eliminate the risk of remote online attacks, while enabling operational and business processes to continue as usual. The gateways physically permit network traffic to flow from OT networks to IT/corporate networks, without the possibility of any traffic flowing back into the OT network.

EA: How does your platform extend unidirectional control into a flexible security solution for ICS?
LF: Waterfall's Unidirectional Gateways are a combination of hardware and software. Unidirectional Gateway hardware consists of a fiber-optic transmitter unit coupled to a receiver with a short piece of fiber. Unlike standard fiber-optic equipment, the transmitter has no receiver, and the receiver physically has no transmitter on the circuit board. The equipment is physically able to send information only one way – out of the industrial network. Unidirectional Gateway software replicates servers and emulates devices to offer the customer off-the-shelf solutions for the most popular industrial software used in the market. The software can also replicate many IT solutions to fit a customer's requirements

for complete protection of industrial networks from remote cyberattacks. Further flexibility can be seen in our DIN Rail form factor, and the myriad of ways customers can configure our modular, rack-mount systems.

EA: Do you see industrial engineers paying more attention recently to cyber security?
LF: Given the frequency of cyberattacks over the last few years, I don't believe engineers in industrial environments have any choice but to pay attention. And there is more interest, however the question is this: Do we pay enough attention to the real difference in approaching OT networks that control physical assets? Many practitioners still approach OT cybersecurity with an IT-based tool set, which, unfortunately, can lead to dire consequences. When OT physical assets are at stake, we need solutions that provide physical barriers against attacks at OT network perimeters to eliminate the possibility of any attack getting through.

EA: How easy (or hard) is it for existing IT security solutions to be adjusted or extended to deal with OT threats?
LF: It's not a question if it's easy or hard – the question is whether IT-based security can succeed at all in eliminating the possibility of an online cyberattack from compromising physical operations via the industrial control network. IT security solutions are all software. All software has vulnerabilities, opening the possibilities to be hacked, and OT networks cannot afford any possibility of being hacked. It is simply not possible to adjust or extend an IT software security solution to adequately protect an OT network from cyberattacks originating from external networks. Take firewalls, for example. Long the standard for first-line defense across IT networks, firewalls are no challenge for modern cybercriminals. Firewalls are porous by nature, meaning they are designed to allow for bi-directional data flows, allowing hackers to easily hitch a ride on a seemingly legitimate incoming message that passes through the IT firewall, which is then used to launch malware inside the IT network to steal business or other data. Now, imagine that same intrusion and the potential impact when hackers breach a firewall to reach an ICS. Firewalls, IDS, and other IT-based solutions clearly have a role to secure corporate networks. They cannot, however, be the sole barrier between a cybercriminal and an ICS.

EA: Do you think the nightmare scenarios so popular in the media regarding OT infrastructure attacks on power systems or nuclear infrastructure are possible?
LF: We've already seen real examples of such attacks, such as the one that shut down the power to a quarter of a million Ukrainians in 2015. The increasing use of ransomware proves that cyber extortion is profitable, demonstrating to cybercriminals that they can improve their fortunes by getting a hold of physical assets until payment is made. Fortunately, in many countries, industrial plants containing nuclear or other critical infrastructure must already be protected by unidirectional gateway technology, which prevents remote attacks from entering an ICS network. In most countries though, far too many critical infrastructure facilities rely on IT-based solutions that can always be breached. Despite the existence of physical cyber protection provided by Unidirectional Gateways, too many facilities and other businesses, like manufacturing and transportation systems, are protected by IT-class solutions leaving them in dire danger of cyber attack.

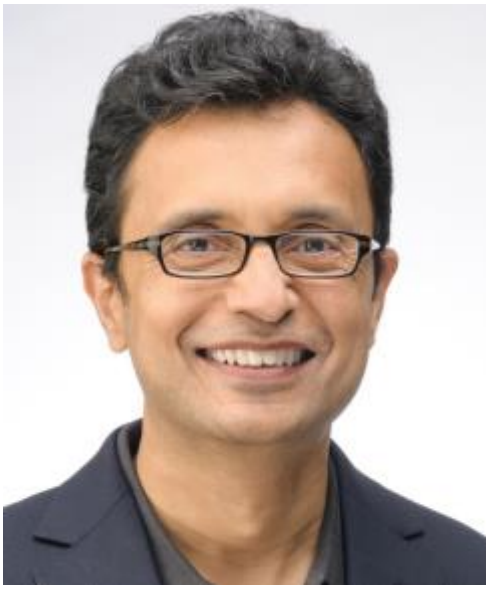

AI-based Solutions for Real-Time Threat Detection and Response

Cyber security AI platform delivers coverage from users to IoT and data centers to the cloud to save time and money

Hitesh Sheth, President and CEO of Vectra Networks

When you work in a security operations center, *time* matters. For this reason, threat-hunting platform designers must follow whatever path is necessary to improve the real-time efficiency of support for the security analyst and threat hunter. It goes without saying that automation must be at the base of this design, but in addition, the use of advanced machine-learning algorithms to detect, triage, and correlate cyber security attacks in enterprise networks can be a powerful means for rapid risk reduction as well. We recently connected with Hitesh Sheth, President and CEO of Vectra Networks, to better understand how all this can be accomplished.

EA: Hitesh, what are some of the challenges of the modern security analyst in detecting threats to the enterprise?
HS: Whenever a new breach is reported, all of us should notice the time lag between the attacker's first intrusion and when their presence was first detected. Typically, it's measured in months. The information security officers employed by government agencies and enterprises are not at fault here; they are well trained, well compensated, and have the best equipment and software at their disposal. So, it is reasonable to ask what is wrong. My belief is that they are overwhelmed by the deluge of security events caused by attempted and successful cyber attacks. These attacks will not decrease — if anything, they will increase. Security operations teams are thus overworked and frequently understaffed. They are dealing with too much noise and low-fidelity signals. These guys are getting burned out doing what is essentially tedious work looking for attackers. This is not scalable.

EA: How important is it for security teams to rely on automation to detect advanced cyber attacks?
HS: Clearly, mere humans can't deal with that flood of security events and sort out the real threats from the pesky nuisances. We must use machine learning and behavioral analysis –

essentially AI – to automate the hunt for threats, and perform triage, correlation and prioritization of those threats inside enterprises. AI and automation augment the human analyst by putting the highest risk threats with rich context at their fingertips so they can act before the attacker causes damage. Once data is breached, there's very little a company or agency can do to recover, regardless of how much money they spend on the effort. Once data is lost, it's lost forever. What we can do, however, is make sure that we detect attacks as they are happening in time to act.

EA: Tell us about how your platform utilizes artificial intelligence to improve attack detection and response.
HS: Vectra AI can spot intruders instantly and tell the security operations team what the attacker is doing with considerable precision. Vectra AI can see what tools the attacker uses and watch those tools morph to improve their concealment. Vectra AI can determine what data the hackers are after, and learn how they plan to move it out of our systems, and we can stop it. We can do all this real-time.

EA: Do you see the biggest risks emerging from IoT, cloud, enterprise data centers, or perhaps all the above?
HS: All the above. Attackers do not see the world in silos. Instead, they look to where the data is, and where the opportunity is. The data lifecycle extends across all these areas, so to find an attacker, you need security visibility everywhere. It is not enough to have visibility in one area, but not in the other. For instance, Vectra has seen attackers hide on IoT devices to launch attack campaigns. We have also seen attackers hide in the virtual infrastructure of enterprise data centers. Cloud is just an extension of those datacenters with the same set of internal problems. You must watch them all.

EA: Hitesh, you have such a wonderful personal background in the network security industry. What advice do you have for young people who might be interested in a career in cyber security?
HS: People aspiring for a career in cybersecurity should view that as an avenue to be a force for good. The impact of cyber attacks is pervasive across our day-to-day lives and is increasingly affecting geo-political situations. This creates an incredible opportunity for innovation for people who are willing to think outside the box.

Distributed Security for Virtual Enterprise

As the perimeter dissolves, the need emerges for virtualized firewall and policy control across hybrid cloud.

Marc Woolward, CTO of vArmour

The original concept of security gateway in the 1990's fit nicely with Internet access. That is, every enterprise was connecting their diverse LAN to the Internet, so placing a firewall there made perfect sense. This gateway concept has dissolved amidst the complexity of remote access, telework, third-party contracts, outsourcing and offshoring, mobile device use, and on and on. Nevertheless, the requirement to manage policy and enforce mandatory controls across this increasingly virtual environment has not waned. Some vendors – like vArmour – detected this trend years ago, and began building effective solutions for enterprise and service providers. Marc Woolward, CTO of vArmour sat down with us recently to help us understand how these important industry trends have been realized in distributed, virtualized cyber security systems.

EA: What is the biggest security challenge for enterprise customers who are moving to hybrid cloud environments?
MW: First, it's refreshing to see so many enterprise teams adopting a hybrid cloud approach. The advantages of using virtual services and infrastructure are becoming obvious, and our team works with customers every day who are aggressively shifting in this direction. The security challenge is essentially the same as one would find for any architectural change. Experts must identify risks, prioritize them, and then implement cost-effective security controls to reduce that risk. The good news is that these steps are simplified when you are dealing with virtual systems. For example, when our vArmour solution is integrated into a cloud infrastructure, the deployment is light, virtual, and involves no new hardware or the need to shoe-horn a gateway into a naturally flat, scale-out network architecture. This is also true for most cloud security solutions – the deployment is simpler. The other issue that has emerged is the need to understand your applications to secure them in potentially public, multi-tenancy environments. To implement security policies, you need to understand your application's dependencies, but

also security best practices. I would say that this is the area that operators need to address, but the good news is that this can be done in an automated manner, and in a way that will improve overall application security.

EA: Do you see virtualization as a security challenge or as a security solution?
MW: All technologies, including virtualization, introduce new security challenges, particularly when you constrain your thinking to legacy approaches involving appliances. This means that network functions such as distributed policy management, which are required in a cloud environment, certainty must be selected and implemented. So, there are challenges, but virtualization also provides some tools to implement distributed controls executing dynamically in virtualized namespaces, along the lines of virtual network functions (VNF). The larger context is that existing perimeter-based solutions are not working. By adopting focus in the data center or network on software-defined virtualization, the overall risk will drop accordingly. In this sense, you could say that technologies such as virtualization and cloud are important parts of an overall security solution for enterprise. With these advances, enterprise teams gain access to application-aware monitoring and reporting, cyber deception, micro-segmentation, and other software-based advantages that do not come with traditional perimeter solutions. Stepping beyond host-virtualization to OS-virtualization, with its containers, Docker and the like, it is also important to ensure that your controls can address micro-service-level security, because a larger attack surface gets exposed to the network via APIs. Fortunately, we have found that the same sets of distributed systems principles apply equally to securing containers as to VMs.

EA: How do enterprise customers keep track of all the policy enforcement points scattered across cloud workloads?
MW: That's the essence of what we help our customers ensure when they move workloads to cloud. Some people refer to this as orchestration, and you are correct that policy enforcement points will become scattered. Remember, however, that existing policy enforcement on a global perimeter is basically distributed, albeit within the same logical perimeter. The difference in cloud is that the workloads will be hosted on a variety of underlying infrastructure environments, which is why it is usually called hybrid cloud. Keeping track of all this can only be done by automating the orchestration task, and providing tools for ensuring consistency in policy across the virtual edge. Fortunately, deploying the security controls at the edge, adjacent to each workload or application, not only makes security stronger, but also eliminates many of the path computation issues you will find with traditional networks. You can be sure that the policy enforcement point adjacent to the workload is responsible for its security. It is the job of a distributed security system like vArmour's to abstract away environmental differences and topologies across hybrid clouds. That job is made easier with a model of deploying security at the edge, and thus not needing to manage complex service chains.

EA: Is mobility an important consideration for enterprise organizations moving to cloud?
MW: Mobility and cloud go hand in hand. While mobile devices have certainly come a long way in terms of performance and capability, the real power of having a smart phone, tablet, or even IoT device is the cloud interface and the amazing content, visibility, and unlimited

networking potential that come with virtually hosted infrastructure. This implies that the security solutions must be coordinated. You cannot do one without the other. From an operator perspective, with cloud you are now operating in highly dynamic, public multi-tenancy environments so you need to understand your application and your threat model. There are tools emerging to automate the computation of application requirements with security control. At vArmour, we think this is incredibly important.

EA: What are some of the big cyber security threats you see coming in the next few years?
MW: We have recently seen advanced nation state attack tools and methods find their way into the hands of for-profit hacker groups. This escalation in capability, partly enabled by source code theft, but also by the development techniques which allow rapid reuse, represents a change to the threat model for many organizations. To me, it further reinforces the need to implement segments within enterprise's networks to create partitions that cannot be penetrated by advanced attacks on common software functions from web to file sharing. Now, automation and the increase in connectivity come at a cost. Container technology and cloud orchestration systems, for example, expose a whole new attack surface from all those APIs and services that communicate with each other. If you are building a cloud, you must ensure that you understand how to secure those interfaces because they provide a new vector for attackers. My view is that this new risk more than offsets the potential benefits. I also worry about protocols that are necessary for the functioning of the Internet that were not designed for hostile environments supporting protocols like DNS and BGP. If your threat model includes nation state actors, then advances in computing models, specifically quantum, will have an impact on efficacy of today's encryption algorithms. There are suggestions that cryptographic transport meshes are the solution for everything. First, that's only as strong as your implementation, but it also obscures what is happening once an attacker has gained access and is potentially ineffective against attackers with access to advanced computing resources. Once again, the case for segmentation of infrastructure, along strong cryptographic authentication and protection of data at rest, provides a balanced mitigation.

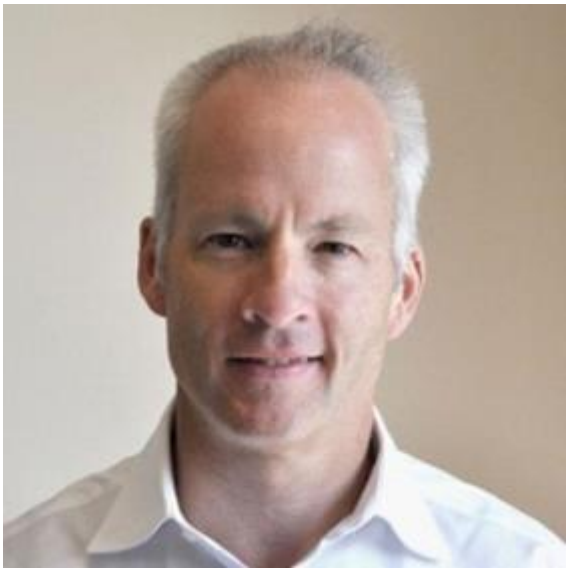

Enterprise Use of Threat Intelligence Exchanges

Real-time cyber security posture depends on accurate and current intelligence obtained from trusted groups.

Paul Kurtz, CEO of TruSTAR

Threat intelligence sharing is the most common security mitigation tool cited among executives, decision makers, and government officials. However, it's also one of the most poorly understood controls in the modern security enterprise arsenal. Modern CISOs know that threat intelligence must be carefully managed and must come from sources that can be trusted to offer reasonable, accurate, and meaningful information. Without these attributes, an enterprise team can waste time on bad information embedded in useless threat intelligence feeds without any actionable context. Paul Kurtz, co-founder and CEO of TruSTAR, is an expert in this area and spent some time helping us understand how this is best achieved in the enterprise.

EA: What is the best way for enterprise security teams to share information?
PK: We have learned that while most organizations want to share, they are not ready to do so. Many organizations struggle to map their internal threat landscape, which makes it difficult to decide what to share. After some trial and error, we understand there are three requirements which lead to effective sharing. First, companies must be able to seamlessly correlate events *inside* their organization. Often teams can't reconcile current events with past events to see how events inside a company are related. Second, they need to be able to operationalize threat data from outside parties, such as ISACs or proprietary threat feeds. Third, companies want to understand their return on investment before exchanging threat data, meaning they would like to see how their events relate to others before engaging in active sharing. When companies do decide to share, they still want anonymity and the ability to redact information on the fly. Organizations like the Retail Cyber Intelligence Sharing Center (R-CISC) or Columbus Collaboratory have successfully enabled sharing while protecting the identities of their members. Through these steps, organizations are engaging with each other and receiving real-time threat insights from other companies. In fact, some of our partners estimate that threat intelligence sharing has helped them reduce fraud investigations by as much as 1,200 days.

EA: Is automation a requirement now for good threat sharing?
PK: Great question! Of course, it depends on what you mean by automation. The reality is that we are still ways off from machine-to-machine event sharing, even though we have the STIX/TAXII standards in place. The more important automation question is how threat intelligence platforms (TIPs) seamlessly fit into a company's' workflow. For example, you need to be able to merge data from email, SIEMs, orchestration platforms, and ticketing systems inside companies. In the absence of integrations, it becomes a very manual process and the return on investment drops significantly. I do believe we will see the day that we can use protocols like STIX, but most companies are not yet close to that objective. TruSTAR recently rolled out an automated email ingest capability that has been getting a lot of engagement from users. We found that many companies who pay for ISAC/ISAO memberships could not tap the value from their industry sharing groups because they received indicators via unstructured data formats like email. Once we added the email ingest capability, we could relieve security operators of the mundane task of manually inputting data from email into their TIP. We also have customers that use a combination of Splunk and ServiceNow to engage in sharing. It is proving to be a very powerful combination.

EA: How does an enterprise go about developing or joining a trusted sharing group?
PK: Many such groups already exist. If you are part of the financial sector, for example, then sharing groups already exist – and our platform can and does provide support to such organizations. So, developing sharing relationships often occurs first, and then implementing technical support for threat exchange functions are done next. If you have the opportunity, we recommend developing a sharing community in conjunction with the use of a TIP like TruSTAR for optimal process definition. When it comes to joining sharing groups, TruSTAR recommends a crawl, walk, run approach. First, you should get your house in order. Enable company operators to understand how their event data correlates internally and with other companies' incidents before exchanging data with others. Second, you should operationalize existing relationships. Operationalize information from other threat feeds or relationships based on your events. This could come from information sharing groups like ISACs/ISAOs or sharing groups of your own making. Do not drive operators to run time-consuming queries of other sites hunting for data. For example, enable seamless use of data from an Information Sharing and Analysis Center (ISAC) or an independent provider such as CrowdStrike. Third, you should scale intelligence exchange into your SOC. Exchange data with other companies in the network based on relevant correlation and without relying on a third party for attribution protection.

EA: Can smaller companies benefit from threat intelligence sharing?
PK: There are no size or scope issues that might limit an organization's ability to share and benefit from a threat exchange. We find that larger companies are starting to fold in their supply chains to their sharing processes. Also, we find that MSSPs are adopting our platform to serve smaller customers. This is a terrific use-case as some smaller companies don't have an in-house security staff and rely on others to assist. MSSPs leverage TruSTAR to capitalize on the overall network effect of our exchange model, bringing greater insight to their customers than if they were to operate independently.

EA: Tell us a bit more about your specific platform and how it supports sharing objectives.
PK: TruSTAR is designed to operationalize threat feeds and to support the listservs in ISAC and ISAO groups. We've designed our platform to support event correlation so that data can be translated into useful intelligence with actionable context. Our IOC exchange functionality supports the creation of sharing amongst a community of relevant peers, and we auto-redact sensitive information so that privacy can be preserved with minimal risk. Overall, we emphasize accuracy, speed, confidentiality, and flexibility for any group of companies that wish to benefit from threat information sharing.

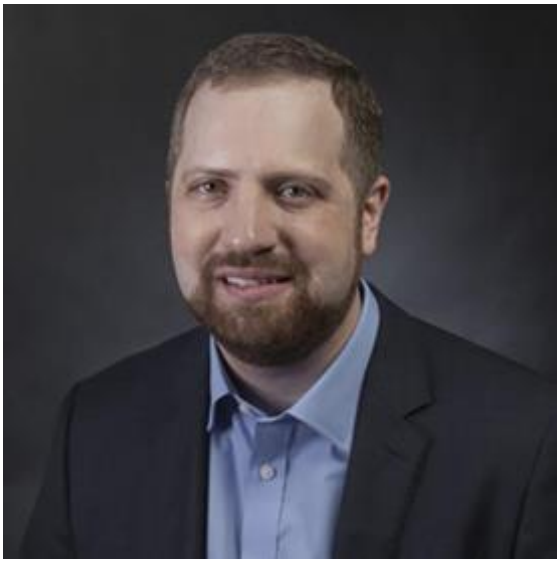

Foundational Cyber Controls for Security and Compliance

Risks emerge in the enterprise from insufficient attention to file integrity, configuration control, and other administrative tasks.

David Meltzer, CTO of Tripwire

So much of modern cyber security involves developing new solutions to threats that are presumed to be on the horizon. Few companies, however, have the experience and legacy to include current solutions that deal with exactly the types of security issues that enterprise customers are dealing with today. Tripwire is just such a company, having been a leader in our industry for as long as any other firm, and possessing of experience and expertise developed through many years of customer security support. David Meltzer, CTO of Tripwire, shared his views with us recently on the advances his team has made in practical management of logs, vulnerabilities, files, and other modern enterprise assets.

EA: How does configuration management factor into the enterprise security ecosystem?
DM: Misconfigurations, many of them easy to correct, have been the underlying reason for many successful breaches. Secure configuration management (SCM) is the control that assures systems are set up and maintained in a way that minimizes risk while still providing the essential business function of the system. Maintaining configurations is so vital to an organization's data integrity that just about every security framework and compliance regulation related to security calls for SCM. While SCM can seem simple in a small organization, it's quite complicated for enterprises that operate larger, more complex technology environments consisting of numerous systems, asset owners and applications - all with differing configuration states and business requirements. For this reason, enterprises would benefit from technology that automates the assessment, monitoring, and management of configurations across all systems to ensure ongoing security and compliance.

EA: How about file integrity monitoring – how do CISO teams provide for this important control?
DM: These days file integrity monitoring (FIM) might be more accurately described as "system integrity monitoring" – which is a fundamental and foundational security control

because it answers the key question: Are systems still in a secure, trusted state, and if not, what changed? What we commonly refer to as FIM has evolved quite a bit over the years. I think of it now as a broader process, not just about monitoring changes for *files* but also the integrity of *registries, databases, and applications*. FIM has also evolved to go beyond just getting visibility of the changes. A good FIM or system integrity monitoring program should also then be able to sort through and prioritize those changes to help you build an actionable workflow for addressing them. For example, is a change introducing risk or non-compliance? Does that change go outside the established organizational or regulatory guidelines?

EA: Do most enterprise teams deal with vulnerabilities in a proper manner? Do they need automated support?
DM: On-going exploits of known vulnerabilities show that vulnerability management (VM) is still a challenge for many organizations. Most large organizations have some form of VM in place, but generally there's a lot of opportunity for improvement. We see a lot of VM programs demand time and manual effort from their teams, so it's a matter of VM programs maturing and incorporating more automation. VM can be hard to tackle when dealing with data overload and relying on slow and error-prone manual analysis. Some specific questions to answer when maturing your VM program include the following: Are you scanning everything that needs to be scanned? Where are you deploying your scan engines around your network? Are you using credentialed scans? How quickly are you able to remediate? How efficiently and accurately are you able to prioritize risks? Do you have the right metrics? How many of your assets you are scanning? What is the effectiveness of remediation? How well is vulnerability information being communicated? Are the asset owners aware of the findings? Are there executive dashboards available for upper management? Is the SOC getting this information? Would your IT service management team can benefit from knowing your vulnerability state? These are numerous questions, but they are all vital to proper VM.

EA: Are security logs managed properly in the enterprise?
DM: Sifting through mountains of log and event data can get overwhelming. In today's environment, what you really need is log *intelligence*, with security analytics and forensics for rapid response. Although almost every organization we work with has some log management system in place, there's often a lack of actionable information coming out of those systems to help reduce risk or prevent breaches. Although just collecting the logs may be a valuable way to prove compliance, organizations should explore use cases that will help reduce risk and enable them to proactively identify potential issues.

EA: What are some of the future trends you see in enterprise security and compliance platforms and solutions?
DM: Maybe this isn't so future, as it's happening now, but are seeing massive adoption of three themes related to cloud: use of public cloud, adoption of DevOps, and the use of containerization in application development. For most large enterprise, their future will be hybrid – environments combining physical servers, virtualization, and both public and private cloud. Visibility and the implementation of a consistent set of security controls across these systems will be needed to maintain strong security postures in this new mixed

environment. More organizations will continue to adopt DevOps practices, and security teams will need to try to keep up with new processes and technologies that introduce different kinds of risks and challenges. Containerization is an especially interesting trend to follow in terms of security. Maintaining visibility of containers and their contents can be challenging, as they tend to be numerous and change often. Security teams will need to keep up with their DevOps teams to implement proper security controls on the contents inside those containers. There's been good progress in this area but we'll see this continue to evolve.

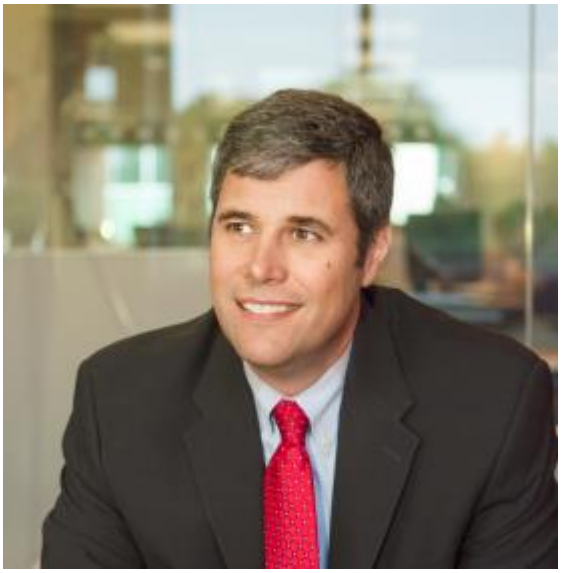

Embedding Cyber Security into Utility Network Services

Attacks at the lower layers of the network stack are generally made against poorly maintained networks.

Bruce Flitcroft, CEO of TenFour

Just as energy services can be procured as a utility, so can underlying network services for the enterprise. Through partnership with the best ISPs and network vendors in the world, companies such as TenFour have been able to construct utility services that integrate the lower layers of the protocol stack into a combined set of services that simplify network operation for the enterprise. This has useful implications for cyber security, simply because standard protection components can be embedded into the utility service that federate and export security indictors, intelligence, and action to the upper application levels. Bruce Flitcroft, CEO of TenFour, made himself available recently to share his insights into how security is supported in this type of network arrangement.

EA: What is meant by utility infrastructure services?
BF: What we've pioneered at TenFour, which many of your readers might recognize by our former name Alliant Technologies, is the design and delivery of a set of standard IT utility infrastructure components into an agile and reliable on-demand network solution. We've taken all the core IT infrastructure that was previously "uncloudable"—from routers, switches and firewalls to phones, WiFi, cameras and IoT devices—and deliver them as a utility service. We've even included all the bandwidth and circuits. As you know, we also embed and integrate security into this concept. The result is that our customers let us take care of the lower layers of the protocol stack so that they can focus on furthering their business agenda and digital innovation while protecting the business.

EA: Do you see many threats hitting enterprise customers at the lower network layers?
BF: Unfortunately, the answer is yes. Before we get there, the customer is getting hit because their surface attack area is enormous and irregular. We use a reference architecture design with smaller and more simplified surface attack areas. As a result, we

see attacks decrease since there are easier targets elsewhere. The challenge we take on is to make sure that these threats do not create serious problems for our customers. We use standard components to build sensible security protections for network layers 4 and below, and we export the alarms, logs, and notifications we receive through our service interface to customer security systems such as security analytic platforms and SIEMs.

EA: How are applications better protected by using a more secure underlying network base?
BF: It was probably correct to say that the earliest original security attacks clearly targeted the lower layers of the network stack. We all remember those early TCP/IP packet attacks that hackers liked to launch in the Nineties. Today, however, the biggest security challenges seem to exist at the higher levels, usually targeting applications and users. Given that a house is only as strong as its foundation, the TenFour team recommends integrating security solutions into the underlying utility to free up the security team to focus on attacks to applications. Every network requires a multifaceted security plan that should be diligently maintained so there are no cracks in the foundation.

EA: How do utility services deal with DDOS attacks?
BF: We approach the problem using the best standard solutions from service providers to divert traffic and ensure proper scrubbing. The challenge, as mentioned above, is that many DDOS attacks are moving up the stack and beginning to target applications. This requires more tailored solutions based on the specifics of the application. Our utility service is designed to support this activity by ensuring solid network controls.

EA: What are your predictions for the coming years in this type of utility network protection?
BF: The TenFour team believes that utility solutions will increase in relevance and importance across the entire IT industry. More and more security features, such as log management, access controls, intrusion detection and firewalling, are just going to be a requirement of the standard service and not sold as standalone elements. TenFour has taken this approach by embedding network security as a core service of its IT infrastructure utility. As standard, automated components can be used to create best-in-class networks for enterprise, it makes perfect sense to move toward this more efficient approach. Accordingly, we believe that more enterprise teams will come to recognize and rely on utility-based network protections. As attacks move up the stack, it is a good idea to deal with the lower layer attacks in the most standard manner possible. Utility security solutions work that way.

Hacker Powered Vulnerability Orchestration

Advanced risk reporting and penetration testing benefit from crowd-sourced analytic platforms

Jay Kaplan, CEO of Synack

Early bug bounty programs were little more than websites with contact names and vague promises of fair compensation. These programs quickly evolved into better organized solutions to the challenge of nurturing relationships with security researchers and enterprise organizations. Several themes have dominated the more successful solutions in this area. Platforms have tended to be more powerful than point solutions or processes; crowdsourcing has tended to be more powerful than individual research activity; and well-defined researcher compensation has tended to be preferred over ad hoc decisions made after a vulnerability has been reported. Furthermore, the community is finally realizing the potential for crowdsourcing to replace the highly-commoditized penetration testing space. Jay Kaplan, CEO of Synack, has been a pioneer in this important aspect of enterprise security. We asked Jay to share his unique insights into the evolution of crowd-sourced vulnerability orchestration.

EA: Jay, tell us about the services you offer at Synack.
JK: Synack is reinventing the way organizations conduct security testing. When my co-founder and I were at the NSA, we saw that red teams were scarce and that static pen tests were not providing organizations with the data they needed to harden against attack. We set out to change that. We founded Synack with the goal of providing a scalable, offensive approach to defense. Today, we offer the hacker-powered security platform as a managed service. Through our platform, we provide on-demand, scalable crowdsourced pen tests to F500 companies and government agencies. In real time, our clients receive analytics and reports on what, when, and how our global crowd of vetted ethical hackers is testing their assets. All this data is filtered through our internal team for quality and impact so that security teams can focus on what matters most: Making their organizations more secure. Our team works with clients to measure, manage, and improve their attacker resilience over time through our continuous testing model. We want to make it increasingly difficult for the adversary to find and exploit vulnerabilities in their systems. We fully manage,

incentivize, and vet our global team of ethical hackers, the Synack Red Team (SRT), to provide our clients with the best talent in the world without any hassle. The SRT brings an unparalleled diversity in perspectives and insights and they utilize the latest attacker tools, techniques, and procedures to mimic the activity of real world malicious hackers to detect exploitable vulnerabilities. They demonstrate deep specialization in web and mobile application security testing, network and infrastructure security, connected IoT device and embedded device hacking, and physical security/special projects. Because of our man-and-machine approach, our solution can easily scale to the size of modern attack surfaces. Our security-as-a-service model deploys within 24 hours.

EA: Can smaller companies begin to benefit from these types of bug bounty-related services?
JK: Cyber threats are everywhere. While smaller companies may not hold as many valuable resources, intelligence, or power as a larger company, the results of an attack can be devastating. When a small company gets breached, they still lose revenue and hard-earned value. So, smaller companies can and should utilize these services and benefit from them. While we've started with government agencies and large F500-type companies, we've started to see substantial interest from smaller companies as well. The crowdsourced, bug bounty model might help smaller companies more, because they are less likely to have enough resources in-house to conduct effective security testing.

EA: How do you ensure that vulnerability investigative tasks don't create damage?
JK: Bug bounties can be open or private. At Synack we believe only in private bug bounties that leverage highly vetted hackers and track their activity to mitigate any risk and capture insights for the customer. Only about 10% of researchers who apply are accepted to join our Synack Red Team. We have a multi-step vetting process that includes a variety of background checks, interviews, and skills assessments. We ensure that the hackers who work on our platform are ethical and trustworthy. Synack's proprietary full-packet capture gateway technology, Launchpoint™, continuously monitors and captures all researcher reconnaissance and pursuit efforts. The assurance and audit log capabilities of our platform provide additional layers of transparency and trust, allowing our clients to take advantage of bounty-driven application testing for even their most sensitive applications and internal environments.

EA: We all recognize that penetration testing and other security testing solutions have been highly commoditized, ineffective, and heavily reliant on the consultants performing the work. How has Synack addressed this problem while still offering the comprehensiveness of a standardized pen test?
JK: Pen testing is a critical tool, but traditional static pen testing is not wholly sufficient. If a vulnerability is common and you can find it on the OWASP top 10 list; a traditional pen test might help you find it. But if the vulnerability isn't a common one, or if you want to have higher confidence in your security and your reports to the board, a traditional pen test will most likely miss the mark. At Synack, we're taking the best of penetration testing, but with a scale, diversity and effectiveness far superior to anything else on the market today. Our Synack Red Team brings an unparalleled diversity in perspectives to mimic the activity of real world malicious hackers to detect exploitable vulnerabilities. Synack's Coverage Analytics backs it up with real-time data on researcher participation, total active hours of

testing efforts, and the breakdown of Synack Red Team traffic activity on a customer's assets, classified by attempted attack techniques. Coverage Analytics provides our customers with insight and visibility into the testing comprehensiveness of our researchers. Utilizing this feature is critical in ensuring that an organization is becoming more and more resilient to attack in cyberspace.

EA: Do you see Bitcoin as a future preferred means for compensating researchers?
JK: You can't conduct "daily life" transactions, like buying food or paying rent, with bitcoins. Until that day, BTC will remain the currency of illicit transactions or currency conversion tool where banking is hard. People want money they can use now, and given that our crowd scales globally, the needs are diverse. So, for now, today's currency is the best form of payment. Given bitcoins represent currency for the digital economy, one could argue that some hackers might prefer it. However, it will be a long time before the adoption of BTC will scale across to the physical world, infiltrate different locales and become useful in the consumption of goods and services.

EA: Can you broadly outline some of the more interesting vulnerabilities you've seen reported from your Synack Red Team?
JK: In one example, there was a logic flaw that resulted in an authentication bypass where the attacker could login to a website with a valid account and through a redirect to SSO, they could gain full admin privileges. The attacker could intercept a redirect to directly access the dashboard that granted full administrator access. In this example, the dashboard assumed that if you're logged in on this page, you should be granted automatic admin privileges. Being a logic flaw, this vulnerability is next to impossible for automation techniques to find. Standard methods didn't catch it, because automation doesn't see to this depth and penetration testers just overlooked it. It happens, and it's why you need lots of eyes looking at the same problem. What's maybe even more interesting than the vulnerabilities themselves is comparing the time it takes to find a given vulnerability between an internal security team and our Synack Red Team. In one case, for the discovery of the same vulnerability, our Synack Red Team made the discovery in a matter of hours versus 3+ weeks of the client's internal teams. Having a hacker mindset helps to accelerate the discovery of vulnerabilities like this. In fact, in more than 75% of environments with competing solutions in place, our Synack Red Team discovers an unknown vulnerability of higher severity (CVSS 7+) within 24 hours.

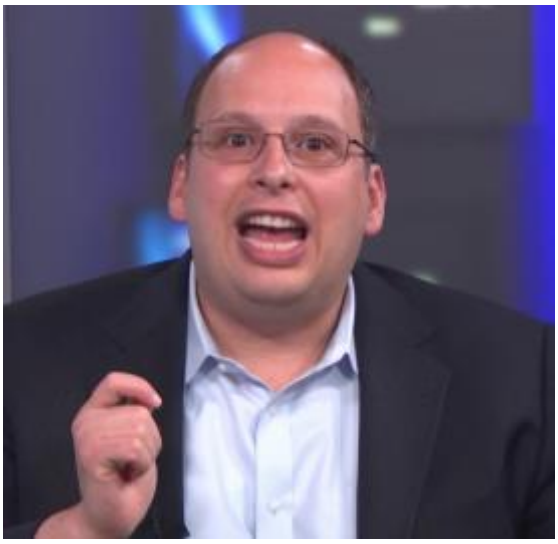

Next-Generation Cyber Security Solutions

Advancing new and innovative protections for modern enterprise cloud, networks, and applications

Hugh Thompson, CTO of Symantec

Few companies reach iconic status in our industry, but Symantec has earned that position as one of the great leaders in the cyber security community. With the recent merger of Blue Coat and Symantec, a new management team has brought new energy and capabilities to the company – and it's hard not to get excited about their fine security services. Combining the world-class proxy solutions pioneered at Blue Coat with the endpoint, network, and infrastructure capabilities of legacy Symantec results in a powerful new resource for cyber defenders. Hugh Thompson, CTO of Symantec sat down with us recently to share his views on industry and threat trends.

EA: Hugh, how complex was the process to integrate Blue Coat, Symantec, and other acquired entities one combined organization?
HT: Though the process of integrating various technologies can be complex, there were terrific synergies and almost no product overlaps between the Blue Coat and Symantec technologies. This has also been the case with our other recent acquisitions. We look to acquire technologies that are additive to our product platform, and we have the singular talent and resources to get the job done quickly relative to most any other vendor in the industry, all of it driven by the needs of our customers. We also have another big advantage in product integration. Early on, Blue Coat had a philosophy of an open architecture that made our products highly extensible and easier to integrate with third party solutions. This open ecosystem greatly accelerated our ability to integrate the Blue Coat and Symantec product sets and this philosophy has now been adopted by Symantec as well. As to examples of the success we've had integrating the two product sets, one early milestone was the combination of the Blue Coat and Symantec threat intelligence feeds, which created a differentiated lens into the threat landscape and which was completed just weeks after the closing of the acquisition. To fully understand the power of this integrated threat

intelligence you don't need to look any further than the significant value this combined data now brings to our artificial intelligence capabilities in blocking an *additional* 3.2 million attacks every day. We likewise continue to expect more benefit as our Integrated Cyber Defense Platform ingests this threat intelligence and combines it with robust integrated functionality across our user, web, information, and messaging solutions. Another more recent example of our rapid integration is the feature parity we achieved between the Blue Coat and Symantec network products delivered through appliances and their corollary in the cloud. We now have leading software elements across all relevant deployment methods – appliances, virtual appliances and pure cloud – giving our customers the flexibility of running our products wherever and however they'd like. Again, we were able do this in record time because of our open architecture. It's also important to note what we accomplished in combining Blue Coat's cloud proxy and CASB, and Symantec's data loss prevention and multi-factor authentication into one complete cloud offering that addresses new risks. For enterprises to take advantage of the cloud, security solutions must reduce the risks inherent in the cloud generation. With our integrated cloud offering, we now possess all these major components. The industrial logic around the integrated Symantec and Blue Coat products as well as the integrations of other companies we've acquired is stronger than ever. And we continue to move faster than ever to combine leading product integrations and innovation with unmatched scale. All of it gaining us a significant competitive advantage.

EA: What technology trends are driving security solution designs at Symantec?
HT: With enterprises shifting IT workloads over to the cloud, embracing cloud technologies in an unprecedented manner, more pressure than ever is being put on the endpoint. Enterprise networks are not going away but they are rapidly being augmented by cloud-delivered applications and services. This make the endpoint an increasingly important place to protect enterprise users and data. This transformation is the reason for our innovation in areas like CASB, and it's also the reason that endpoint security is fast becoming one of the key drivers for our business. At this moment in time, we're laser focused on delivering a converged endpoint that will reduce risk and lower expenses for our customers. We've always had some of the industry's best endpoint protection technologies and we've very quickly closed the gap with competitors around artificial intelligence and machine-learning-based detection. We've also turned our resources and energy toward Endpoint Detection and Response (EDR) a critical piece of the endpoint puzzle that gets to the remediation and response of malware and cybersecurity events. At the same time, with end users' work and home life becoming harder to separate, we are starting to see an acute need to protect people and their data as they move between being enterprise users and consumers. Utilizing our Norton consumer security products, we're paying a great deal of attention to developing the world's first comprehensive Digital Safety solution for consumers, delivering more value than PC malware protection alone and protecting all aspects of a consumer's digital life, including their information, identities, devices, homes and families. Trends in online consumer activities are mandating that companies like Symantec step up and deliver solutions that take the lead in protecting users and their families whose lives are increasingly dependent on their online transactions and communications. And protecting these users as they move between their personal and working lives also protects the enterprises we serve as well.

EA: Do you see cloud services and virtualization as the new playing field for cyber security technology providers?
HT: Yes, cloud services have created a new playing field, but the play that many of our customers are running as they advance into the cloud often deploys a mix of cloud services and strategic on-premise deployments, either through traditional or virtualized appliances. There are many reasons for this – from data regulations that must be met to direct-to-net traffic that still originates from corporate headquarters – and even though a full cloud deployment model is enticing, the reality is that most enterprises are looking to take it one step at a time with the total security of their organization first and foremost in their minds. This puts Symantec in the right place at the right time as we offer both the cloud services infrastructure and on-premise deployment expertise to make sure our customers get what they need. Our cloud services have become world class and the subscription service go-to-market machine we've built to deliver to services is transforming our business. The future of this game is cloud, but for Symantec the field we play on is dictated by our customers, namely, what they need us to deliver in the deployment model that best secures their users and their business.

EA: What do you see as the future of endpoint security protections? Is signature-based antivirus dead?
HT: As we talked about earlier, with enterprises shifting IT workloads over to the cloud, endpoints have become a critical area for security control. That said, endpoint protection and control will never be just a game of one technique against another – that is, signature-based antivirus versus artificial intelligence or machine learning. No matter what some vendors try to tell their customers, to truly protect and secure the endpoint takes a combined effort of multiple techniques and strategies. That's why we continue to invest and lead the way in artificial intelligence and machine learning *as well as* in signature based antivirus which still has a role to play to stop known malicious activity that AI may miss. It's also why we recently acquired Skycure, a mobile endpoint threat defense technology that utilizes advance techniques which can be applied to iOS and Android but are also extensible to almost any other device at the endpoint. In short, we don't take anything off the table when it comes to protecting the endpoint where we continue to double down on innovation.

EA: As one of our industry's veterans, what observations come to mind with respect to the nature of the cyber threat in the next decade?
HT: What's most obvious to us daily is the ongoing and massive innovation we see in cyber threats both on the defender side, where we're trying to protect against those threats, and on the attacker side, where they originate. We see cyber threats and the innovation around them being integrated into every aspect of society, from voting to shopping to social interactions to geopolitics. There is barely an aspect of human life that hasn't been touched by the potential for malicious cyber activity or by the innovation needed to protect consumers and users from that malicious activity. This being the case we are going to have to continue to grow and bring in diverse expertise from all areas of study. Cyber threats and the security technology that protects us from those threats are moving from self-contained industries to something that touch every single aspect of our lives. In some very

real ways as cyber spreads into all these things, industry veterans like Symantec must always innovate ahead of the ingenuity of attackers to protect our lives and ensure the safety of our planet. As the world's largest cyber security company, we see ourselves as stewards of the world's digital safety—users, organizations, governments, and everything in between. We don't take that stewardship lightly and it's what keeps us pushing ourselves to do better every day.

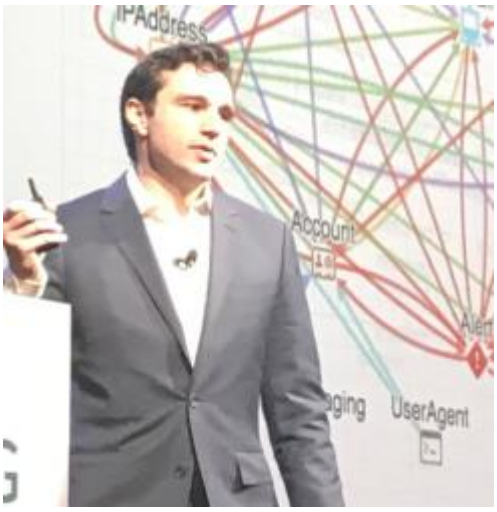

Supporting the Modern Cyber Threat Hunter

Supporting advanced analytics to target, hunt, and disrupt advanced enterprise cyber threats

Ely Kahn, Co-Founder and VP of Business Development of Sqrrl

The rapid transition in the typical SOC from reactive indicator response to proactive hunting of threats is one of the bright spots in enterprise and infrastructure security management in recent years. To support the mission of the modern hunter, world-class tools are required that combine data science with advanced search and visualization techniques to detect threats such as insiders and APTs. Machine learning analytics are a good example of how this combined focus leads to useful platform support. Ely Kahn, Co-Founder and VP of Business Development of Sqrrl sat down with us recently to share his views on the evolving SOC and how his team goes about supporting threat hunters.

EA: What is the mission of the threat hunter?
EK: The mission of a threat hunter is threefold. First and foremost, hunters are focused on finding hidden threats that have evaded detection by their existing cyber defenses. Secondly, and perhaps more subtly, hunters should be focused on taking newly discovered patterns and TTPs and building new ways to automatically detect those patterns. Finally, hunters should be mentors. They are working on the cutting-edge of security and should transfer knowledge to more junior analysts and incident responders.

EA: Can you comment on how analytics has evolved to support modern enterprise security?
EK: We have entered the age of machine learning and Big Data, and the combination of these two trends has triggered the creation of waves of new startups, including Sqrrl, who seek to apply these capabilities to more accurately detect anomalies in vast piles of cybersecurity data. With Big Data technologies, such as Hadoop, massive amounts of data can be processed much more cost effectively and on a timely basis. Machine learning algorithms reduce false positive and false negative alarm rates by continuously adapting to the data and organizational environment.

EA: What background is required for an individual to become an effective threat hunter?
EK: The so-called unicorn threat hunter has skillsets that cut across data science, threat intelligence, network security, endpoint security, incident response, and Big Data (i.e., distributed computing). There are probably less than 100 people in the world with this magical skillset. For this reason, we developed a Threat Hunting Platform that brings together these capabilities and lowers the bar on the skillsets needed for hunting.

EA: Can you help us understand the balance between human skill and automation in the detection of subtle attacks?
EK: We are seeing some vendors talk about "fully automating the hunt." We think this is a fallacy and an example of cyber security marketers creating confusion in many people's minds. If you fully automate a hunt, it is no longer a hunt. It is a SIEM, firewall, or IDS rule. Hunts will always be driven by humans, but a Threat Hunting Platform should simplify the hunt as much as possible through the usage of advanced analytics, visualizations, and playbooks.

EA: What trends are you seeing in the types of threats being detecting in the modern SOC?
EK: SOCs are getting hit from all angles, including "low and slow" attacks seeking to exfiltrate data, and "shock and awe" attacks such as ransomware. The key trend across these different types of cyber attacks is around the commoditization of malware. Adversaries are quickly taking malware, repackaging it, and extending it. No longer do you need to be an expert malware developer to pull off an advanced attack. You just need to know where to go to license it.

Cloud-Managed Hyper-Converged Platform Security

Both hardware and software must be securely integrated to support optimal protections for application hosting.

Michael Beesley, CTO of Skyport Systems

With all the emphasis on cloud, many forget that underlying platforms continue to rely on the usual combination of hardware, software, and the various functional utilities, such as hypervisors, that connect them together. Some common recent trends to shift security responsibility more towards application-level software, ignoring underlying systems, firmware, and operating systems, are ill-advised for certain classes of on-premise enterprise workloads. Simple to use, converged solutions that actively build trust into the underlying compute infrastructure for mission critical virtualized workloads are available and can produce superior security for vulnerable, critical, and exposed applications. Michael Beesley, CTO of Skyport Systems sat down with us to explain the basis for Skyport's platform and how its security-by-design contributes to a converged protection architecture for servers, storage, networking, virtualization, and other functional enterprise needs.

EA: What do you mean by hyperconverged security with respect to your platform?
MB: The concept of hyperconvergence involves cloud-managed systems that combine trusted hardware and software to support critical services and infrastructure. Everyone knows that a system is only as secure as its base, and we believe we provide the optimal support infrastructure for modern hybrid cloud services, and especially for systems that have the highest security requirements. Our SkySecure solution offers an easy to use, cloud managed, virtualized infrastructure that hybrid enterprises can use for their most critical, vulnerable, and exposed on-premise workloads

EA: What is the role of cloud in your platform?
MB: It is an essential component, which is why we always refer to our system as cloud-managed. To provide support for today's cloud infrastructure with systems located across a continuum of public, private, and hybrid data centers, the flexibility of cloud allows system

operators to evolve their architecture and systems operations without having to regress through management changes. The cloud-managed aspects of SkySecure facilitate a turnkey, easy to use, self-service infrastructure with a shared responsibility model between our customers and us. But cloud is not just central to our architecture, it is also central to the use of SkySecure for protecting the ever-increasing set of exposed workloads that are found within hybrid enterprise IT, something that traditional network based security approaches are extremely challenged with.

EA: Are threats best mitigated through a combination of hardware and software controls?
MB: Yes, and this is most obvious for the most important workloads within an enterprise, whether that be a critical control system such as Active Directory, or a vulnerable workload running on a legacy operating system, or an exposed workload talking to the Internet and talking back to the core of the enterprise data center as part of a hybrid application. To properly secure these types of workloads, security teams need to ensure proper configuration, administration, and set-up for the underlying hardware base, firmware, hypervisor, and operating system as well as tight security controls and visibility around the application virtual machine. For these sensitive workloads, all threat vectors up and down the stack must be well covered. The Skyport team focuses on making sure its customers experience this coverage at every layer of the stack.

EA: Do you see CISOs focusing sufficiently on server security in this era of cloud virtualization?
MB: We see a growing number of CISOs demanding the type of hyperconverged protections for hardware and software that we offer at Skyport. This realizes itself in procurement plans, data center security architectures, requests for proposal, and enhanced compliance. This is ultimately a good idea, especially for high priority systems with critical consequences if hacked. As enterprise IT continues to adopt hybrid architecture, there are more and more exposed workloads on-premise. We see CISOs, security teams and infrastructure teams focus more and more on full stack protections for these new threat surfaces with an acknowledgment that the best approach is to run these workloads on converged infrastructure that has full stack security built in.

EA: Any predictions regarding converged security in the coming years?
MB: Obviously, we believe it will become a more important component of security architectures for hybrid cloud infrastructure and modern virtualized data centers. But one place where we expect to see the most intense growth involves security compliance. Demanding higher assurance platforms is both sensible and essential to ensure top-to-bottom protection for sensitive, high-priority, and critical servers.

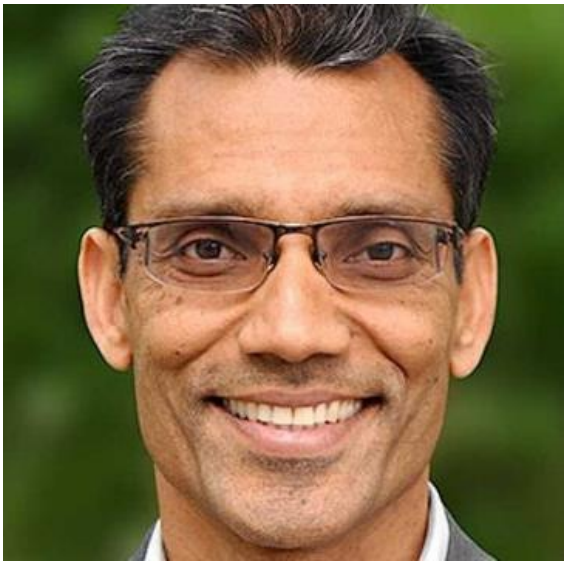

Advanced Cloud Access Security for the Enterprise

Businesses need flexible cloud security solutions that provide visibility and mitigation for hybrid cloud architectures

Rajiv Gupta, CEO of Skyhigh Networks

Every IT team has come to embrace the value of as-a-service offerings in cloud – and this now includes organizations in every vertical industry. Certainly, some organizations are more aggressive in their public cloud use than others, but every organization uses cloud services to an increasing degree. The challenge for cyber security teams is to protect their organizations' confidential data in cloud services to enable confident use of cloud services. This includes preventing confidential data from inappropriately going into unsanctioned cloud services and preventing confidential data from inappropriately leaving sanctioned cloud services. We had the pleasure to interview Rajiv Gupta, CEO of Skyhigh Networks, recently, and he offered some interesting perspectives on cloud access security solutions for hybrid architectures.

EA: What functions does a cloud access security broker solution support?
RG: A cloud access security broker (CASB) must ensure that the use of cloud services by an organization, whether unsanctioned or sanctioned, does not violate the organization's security, privacy, governance, and compliance policies. A CASB platform brings lost functionality visibility, threat protection, compliance, and data security to the cloud. These are the "what" of CASBs. The "how" often determines the value delivered by a CASB. That is, CASB platforms that opt for a cloud-native deployment, and that avoid high-friction architecture like device agents when possible, are more likely to provide the full value of a cloud-native security solution.

EA: How important is visibility of public cloud use to the security team?
RG: Most, if not all, security starts with visibility – you cannot protect what you don't know. In the case of public cloud use, visibility includes knowledge of which cloud service is being used, what is the risk of the cloud service, what activity is being performed in the cloud service by whom, what data is being stored or being created in the cloud service and by whom, and what data is downloaded to which device belonging to which user – essentially,

the holistic context of every data transaction. Visibility also requires analytics. You not only need to know who accessed what data, but through analysis of the user behavior, you need to determine if that user was a rogue insider or a compromised user. Visibility is the start. The real requirement is for insights that lead to action to protect data.

EA: Do you see more enterprise teams converging on a single cloud provider, or are they more often shifting to a hybrid collection of different cloud offerings?
RG: We do not know of a single enterprise team who is converging on a single cloud provider. The reason is simple: There is no one cloud provider who covers the breadth of needs of any enterprise. Enterprises use productivity service providers like Office365 and Gsuite, collaboration service providers like Slack and Spark, CRM service providers like Salesforce and Dynamics, file sync and share service providers such as Box and Dropbox, and many more categories of cloud applications. Of course, many enterprises develop or customize software, and for these, they use hosting IaaS or PaaS service providers like Amazon Web Services and Azure. In fact, many enterprises use multiple service providers for the same function, such as OneDrive and Box for file storage – either because of legacy, transition, purpose, or preferences of their customer, partners, or employees.

EA: What sort of threats do you see in public cloud infrastructure?
RG: The appropriate use of public cloud along with a CASB almost always improves data security. Enterprise-grade cloud service providers typically have better security for their infrastructure and applications than that same application running in an enterprise. With cloud providers specialized in securing their services, enterprises can focus their security investment on the security of their own data under a model of shared security responsibility. Threats in public cloud almost always result from the enterprise not delivering on their part of the shared security responsibility model. Inappropriate use of cloud services can lead to a range of threats including the use of high-risk cloud services, open S3 buckets in Amazon, over-provisioned admin accounts in Salesforce, and storing and disseminating malware. Inappropriate access to cloud services encompasses threats under the umbrella of compromised credentials and rogue insiders.

EA: How do CISOs orchestrate security policies across different public clouds?
RG: To orchestrate a security policy across different cloud services you need to be able to map the security policy to the disparate security controls of each cloud service provider. If a CISO wants to ensure that confidential data is not inappropriately shared, the security team needs to have several capabilities. First, there must be a way to specify that policy, defining what is confidential data and what constitutes inappropriate sharing. There must be a way to map that policy to the different ways data can be shared through each cloud service, which typically offer different actions such as copy, share, invite to collaborate, upload, and download. Finally, they need a consistent platform to get the visibility into the data and to enforce the policy. This mix of cross-CSP administration, visibility, mapping, and control is one of the key capabilities that allows a CASB to enforce cloud security at scale.

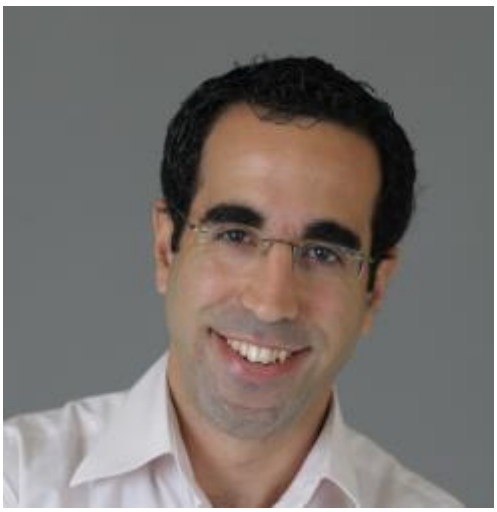

Proactive, Automated Mobile Threat Defense

Using visibility and threat intelligence to reduce cyber risk in the modern operating systems

Adi Sharabani, CEO and Co-Founder of Skycure

Chances are high that you've already made the transition from primary dependency on your PC to primary dependency on your smartphone. This change is exciting because mobile devices are inherently linked to the unlimited computing power, resources, and applications in the cloud. In fact, the synergy between mobility and cloud services helps to explain the dramatic advances predicted by Marc Andreessen years ago when he correctly said that "everyone is going to have a general-purpose computer in their pocket." Now that this mobility vision has been realized, new cyber risks have emerged, and the security industry has responded with a range of advanced solutions based on powerful technologies such as artificial intelligence and machine learning. We recently caught up with Adi Sharabani, CEO of Skycure to learn more about recent advances in protecting mobile devices from cyber-attacks.

EA: Adi, what is your general view of the progression of mobile device risk? Is this the next 'big thing' in cyber security for consumers and business?
AS: Everyone in business today knows that the mobile risk is real, and that working with a world-class mobile security solution provider is imperative. And I'd adjust your question slightly about mobile security being the *next* big thing, because I believe it is a *current* big thing, especially in any business that relies on mobile devices to support individual productivity and business workflow. Such recognition is good, because all evidence shows that mobile threats continue to increase. Skycure reported recently, for example, that for every 20,000 mobile devices, roughly 1,200 have operating system versions with known vulnerabilities. We also reported that organizations with a minimum of 200 mobile devices had at least one malware infection. Judging from these reports, one would have to conclude that mobile device risk is already very high, and still increasing.

EA: Do you see any substantive differences between Apple and Google in terms of mobile operating system security? Does the Apple model, for example, ensure a tighter patching and vulnerability management lifecycle?
AS: The good news is that both Apple *and* Google run more secure software than we've dealt with on our PCs over the past decade, but differences do exist, such as the patching lifecycle. Since Apple controls the ecosystem for its devices and systems, a direct path exists for making security patch updates to devices. In addition, Apple's app vetting process blocks most malicious apps before they can appear on iTunes. Google, in comparison, maintains an open source model that supports multiple OEMs and MSPs. This has obvious business advantages, but it complicates patching and, consequently, security. As much as Apple and Google are doing to improve basic OS security, it will never be sufficient to forego third party protections from the ever-evolving and varied mobile threats. For this reason, despite the security advancements of mobile operating systems, Mobile Threat Defense has become one of the biggest topics in endpoint security. More good news is that IoT operating systems are mostly following the mobile OS model, and even PC operating systems seem to be migrating slowly to the newer paradigm. This explains we did not put "mobile" in our name, and architected our solution to support all the modern operating systems to come.

EA: What's been your team's experience integrating mobile device security with existing or new mobile device management systems in the enterprise?
AS: Mobile device management systems were the initial enterprise security solutions for early adopted mobile infrastructure. While MDM offers essential functionality for business, no one would purport to view it as a security tool. Rather, it provides complementary capabilities such as inventory management and software distribution that the best mobility solution providers must integrate with. At Skycure, we've made it a priority to provide critical risk intelligence that helps the MDM perform its tasks and to integrate with the best available MDM vendor solutions.

EA: To what degree do advanced modern algorithms based on artificial intelligence and machine learning advance the cause of preventing malware risks for mobile users?
AS: The progression from signature-based processing to behavioral approaches to security was a great advance in modern cyber security. We are now seeing a similar progression to machine learning techniques, which take advantage of recent breakthroughs in AI processing efficiency to detect malware more efficiently. Note that the term "machine learning" often means different things to different people, but the Skycure team is excited about the possibilities of this type of approach for improving mobile security.

EA: Adi, what is the most important bit of advice that you'd give to a CISO regarding mobile security risks?
AS: I would advise not to underestimate the mobile threat. Every hacker on the planet knows that individuals and businesses are completely dependent on their mobile devices to function, so attacks on this infrastructure will have a much greater impact than in previous years. This is so obvious, and yet the tendency among many security decision makers is to wait for something terrible to happen before security controls are put in place. This is precisely what I would recommend avoiding.

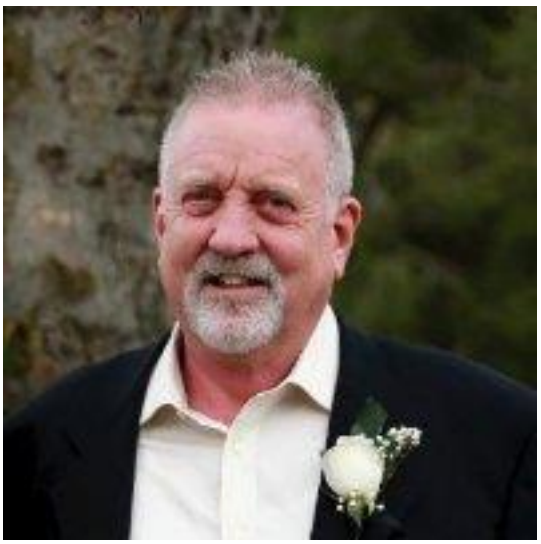

Solutions for Self-Governing, Self-Protecting Data

Using advanced virtualization methods to embed intelligent security controls directly into critical enterprise data

Greg Taylor, CEO of SertintyONE

With the dissolution of the enterprise perimeter, user data is now exposed. Data owners thus have a couple of security options: They can try to build security protections *around* their data, focusing on the environment in which the data resides with the intent to control its users. This technique works fine so long as data remains fixed. If it must be shared or transported between networks, however, the ability to transfer enforcement of policy is complex and the protections can be lost. Alternatively, data owners can build policy enforcement and security protections *into* the data directly – a modern technique that has the great advantage of supporting more transportable data, with greater flexibility and control by data owners. Greg Taylor, CEO of SertintyONE was kind enough to share his insights into how intelligence can be embedded into data to increase security control.

EA: Tell us a bit about your company and how it was formed.
GT: We believe we've assembled an experienced and seasoned executive team that knows how to identify and overcome the unique challenges that we face with a breakthrough technology. Our team includes backgrounds ranging from commercial cyber security to national intelligence defense. Greg Smith joined me in 2010 with the intent of assisting me in relighting two unique online payment systems. Early in the relationship we set out on a path that quickly led to Dan Fischer joining the Company and ultimately, the revelation and discovery of *Smart*Data - an invention and a breakthrough. Our team is in Nashville, Tennessee, which is well-suited to the high technology industry. We'd like to welcome any readers to come and visit!

*EA: Greg, what does your team mean by the term Smart*Data*?*
GT: That's a term we've used to suggest that the data an enterprise needs to govern and protect can no longer be properly secured by the perimeter. Instead, the data requires local, embedded controls, which results in a new type of data – intelligent data – with the ability to learn, act and react to its environment. We refer to it as *Smart*Data. The idea is for strong access controls and

localized encryption to become embedded into the actual data-file. That's about as tightly knit as security can be with any form of structured or unstructured files.

EA: How specifically is virtualization used to embed intelligence into data?
GT: Though not a virtual machine in the sense a full operating system machine, it does represent a small pseudo-code machine that is tailored to the tasks of authentication and data governance. By choosing this approach, *Smart*Data can contain a set of virtual machine modules that permits consistent execution across common platforms, regardless of the data origin. *Smart*Data requires these modules, as data access is impossible without the approval of the embedded code. Virtualization also provides flexibility that is difficult or impossible to replicate using conventional data protection, DRM, or other embedded controls. All decisions regarding authentication, time, location, underlying hardware, or even custom environmental attributes are performed by the embedded virtual machine, which eliminates the need to have a trusted application enforce controls and access. Additionally, conventional encryption and protection are managed by the machine, without exposing keys or requiring external key management. To execute the embedded virtual machine, and subsequently access data, the SertintyONE technology must be present; otherwise, any *Smart*Data object will be inert and appear as a protected blob of binary data.

EA: What are some example use-cases for how customers might utilize your technology to gain increased control of their data as it moves?
GT: There are so many different use-cases, but perhaps one of the most obvious involves an enterprise sharing data with customers, partners, and across business units in a hybrid cloud environment. In such a case, the data traverses so many boundaries that it would be impossible to rely on an external control. Instead, our solution provides enterprise users with the ability to have privacy, security, encryption, and minutia level access control travel with the data as it moves along its path between cooperating participants. The SertintyONE "breakthrough" and the essence of *Smart*Data is the ability to give data an identity, as if it were a user – that is, a processor, an application, a device or a human user. Until now, only a user, application, or machine participates in security processes. Before *Smart*Data, data, unlike users, didn't have the ability to participate and be governed. It isn't that it doesn't have an identity, but rather that it's *not* an identity. It's not a participant. It can't act, or react. The bottom line is we've given architects and developers a tool that enables them to govern or utilize data as if it, the data, were a user.

EA: Do you see your technology supplanting complex measures such as enterprise digital rights management?
GT: Many people draw comparison to enterprise DRM, but we believe that we enable a whole new set of capabilities, above and beyond traditional DRM. That said, we do believe that enterprise users who might have had challenges with the complexities of managing DRM might do well to look at our solution. The same for DLP. Our approach is to integrate our offerings with leading providers, which further eases the technology ramp for most enterprise buyers.

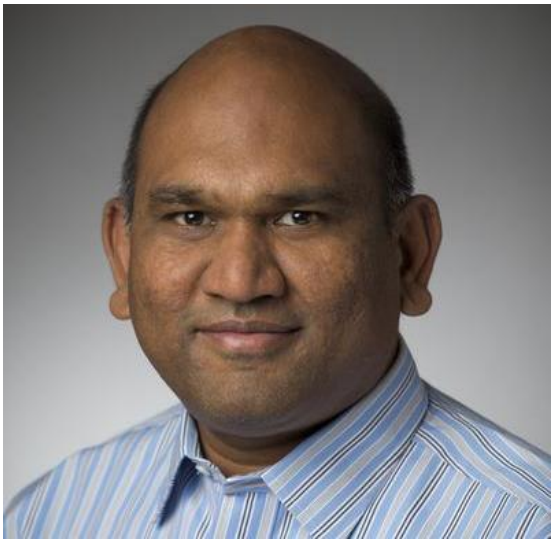

Transformation of Enterprise Cyber Risk Management

Proper prioritization of risk must take cloud, apps, threats, vulnerabilities, assets, and networks into full account.

Srinivas Mukkamala, CEO of RiskSense

Enterprise managers have always had a difficult time determining cyber risk, developing meaningful priorities, and orchestrating mitigation. Vulnerability discovery, for example, is one of the most complex tasks that an enterprise security team deals with on an on-going basis. To strengthen an organization's cyber risk posture, it is essential to not only test for vulnerabilities, but also assess whether vulnerabilities are exploitable and what risks they represent. Meanwhile, to increase the organization's resilience against cyber attacks, you need to have a clear visibility and understating of your attack susceptibility and validate your attack surface. Contextualizing known vulnerability findings with external threat intelligence is especially difficult. Srinivas Mukkamala, CEO of RiskSense, was kind enough to share the details of how his company offers a platform to ease these tasks. Here is what we learned from Srinivas:

EA: What is the underlying risk management methodology your platform supports?
SM: Our platform is developed on the idea that orchestrating the scan output and intelligence from security tools in the enterprise is far too complex a task for human beings. Instead, we believe that an effective platform is required that can integrate with existing and planned cyber security infrastructure to produce clear views of cyber risk.

EA: What is the role of automation in the context of risk management?
SM: Automation is a requirement for any type of scan interpretation and correlation with log and audit output that involves non-trivial size and scope. The modern enterprise security team must have an accurate view of current risk, and the RiskSense platform was designed to automate that task.

EA: As the enterprise attack surface increases, does this impact the enterprise risk process?

SM: We already see the enterprise expending to hybrid cloud infrastructure, so the attack surface has begun to expand rapidly. Mobility, IoT, and related new modern capabilities also increase the likelihood that an attack can occur. As you would guess, this expends the enterprise risk task, simply because it increases the likelihood and consequences of a cyber intrusion or exploit.

EA: How does your platform interact with a typical enterprise IT and security ecosystem?
SM: It is designed to interoperate with the existing and planned systems and tools in an enterprise ecosystem. We've developed connectivity with the most frequently found scanners, and we can consume and interpret output from the security products and tools that will be found in the modern enterprise. We understand that CISO teams have made investments. The RiskSense platform is designed to help optimize this investment.

EA: What are your predictions regarding cyber risk management in the coming years?
SM: It is already emerging as a best practice, so that is consistent with our long-held views and beliefs. It is also a major component of every important cyber security standard, so that is also consistent with our recognition of its importance. Perhaps one area where we would hope the cyber risk management discipline would move is toward greater integration with embedded business unit processes and practices. Some managers still perform risk management in a surface manner, gathering information around existing systems, rather than from within. We expect to see automated platform support bridge this gap.

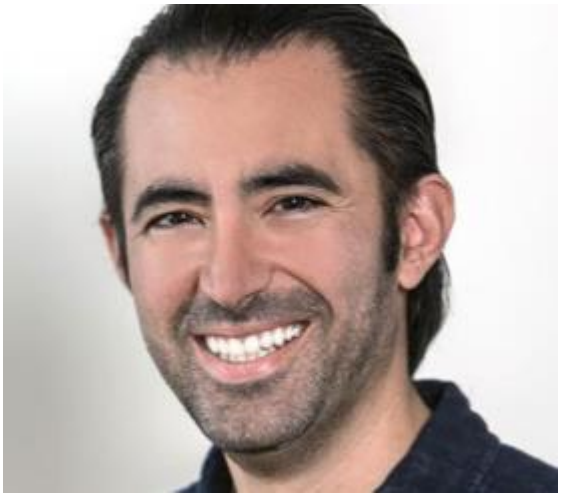

Protecting Brands via Digital Risk Intelligence

Digital threat management is emerging as a primary control in most modern corporate enterprises.

Lou Manousos, CEO of RiskIQ

If there is one constant in cyber security, it is that nothing is constant. In the wake of business digital transformation, threat actors are impacting operations, customer trust, and brands through web, mobile, and social attack vectors. It is simply more convenient for hackers to exploit a business' online presence and exposures in an organization's attack surface, which offer accessible and lucrative targets to commit acts of fraud, misuse, and malicious activity, often duping users to gain access credentials, sensitive and financial information, and system control. Security teams must re-assess their security posture and apply resources, intelligence and controls to mitigate external threats and adversaries. Lou Manousos, CEO of RiskIQ, is an expert in this area, and he spent some time with us to discuss trends in digital threat management and the provision of advanced internet intelligence and response capabilities to support enterprise customers.

EA: What exactly is digital threat management?
LM: Organizations have embraced online mechanisms to enhance product stickiness, customer engagement, and their online ecosystem. Threat actors seize the digital opportunity as well; external threat actors now account for 70% of enterprise data breaches as per the latest Verizon data breach report. Phishing, malvertising, ransomware, rogue mobile apps, web site and app exploits, brand abuse, and fake social posts are all examples of threats that originate outside the firewall. Digital threat management extends visibility and control for organizations across external web, mobile, and social digital channels, and brings that into the fold of security operations. This is more than just threat intelligence data feeds, it is about enabling SOC, red, and blue team resources to gain the insight and automation necessary to efficiently execute tasks that identify, preempt and remediate digital security issues that directly affect their business in an automated way.

EA: Why have so many CISOs not fully addressed digital risks regarding Internet-facing assets such as domains and broader phishing threats?

LM: Adversaries no longer must attack firewalls or maneuver laterally between systems to impact the IT organization and damage business. They can more easily phish with a fake email on online advertisement and website, exploit a susceptible web app or form, create a fake mobile app, or even target a weak affiliate site to feed malware. From our business perspective, most companies only know a small fraction of their Internet-facing assets that can affect their business, and many exposed and exploited assets were outside the purview of IT – and that is what has left companies and CISOs exposed. The sheer volume and respective dynamics associated with all the external assets connected to a business, from company sites and apps to those by service providers, affiliates and adversaries, has outpaced conventional defenses. The general presumption is that current controls, such as those provided by vulnerability scanners, pen testing, next-generation firewalls, and endpoint security will suffice, but *they don't*. For example, relying on point-in-time discovery against only known assets is a false sense of security in the digital world, and reliance solely on reputation services and endpoint updates does not address targeted attacks that are custom and zero-day. The key is intelligence and automation to close gaps in digital defense and to enable CISOs to align with business initiatives while mitigating risks.

EA: How does your team go about collecting and disseminating intelligence on external threats?
LM: We've taken on the investment, technology, and expertise to build out a substantial Internet reconnaissance system comprised of global proxy network, collectors, and scanning technologies to capture data in a variety of ways. We actively scan the entire IPv4 range as virtual users representing different browsers, regions and networks. We collect passive DNS and WHOIS information and more. We actively monitor thousands of mobile app stores and millions of apps. And we have relationships with seven of the leading social networks to actively track posting details. All this data is stored and curated in an elastic warehouse where we apply analytics in the form of correlation models, pattern matching algorithms, data science and research. While the data can be consumed as feeds for some organizations, we have three popular products delivered as SaaS web applications that leverage this data. Our Digital Footprint tool enterprises to understand, monitor, and remediate exposures in their digital attack surface and track their risk rating. Our External Threats tool allows the SOC to identify digital attacks, to triage the issue, and automate response. And our popular PassiveTotal tool enables incident responders and researchers to investigate external threats, adversaries and exploits.

EA: What sort of mitigations can be performed when an organization is under attack?
LM: Certainly, after you identify an external threat, you want your SOC and other security team members to be able to make informed decisions and act. In many cases, our correlation models unearth external threats as they are being weaponized – allowing the defense to preempt attacks. We have automated mitigation tasks across digital channels, and offer an extensive API set for interoperability with custom and popular security systems. For digital threats, including targeted attacks which utilize newly identified phishing and malicious sites, we can update blacklists for blocking within Firewalls and web filtering tools. We can enrich threat data into SIEMs. And we can send our external asset discovery data into GRC and VA tools. In addition, the RiskIQ platform can

dynamically submit phishing, scams and other malicious URLs directly to Google Safe Browsing and Microsoft SmartScreen to block these URLs within 95% of web browsers in a matter of minutes. For digital threats where the organization has little to no direct, immediate control over an external asset, our platform offers automated takedown workflow, and monitoring. This would cover such digital threats as domain, mobile, social, and brand abuses, where it requires contacting an entity for dispute and corrective actions, including their support infrastructure, such as registrars and hosting providers.

EA: Any advice for organizations in this area to get started? Can small companies afford to purchase digital threat management?
LM: Small to medium enterprises can and should perform an assessment of risk with regards to their potential exposure to digital threats. We have packaged our products to even accommodate organizations with limited means and resources. In fact, RiskIQ has a Community Edition of our tools that gives entry level access to our information and tools at no cost. This is for our PassiveTotal threat investigation tool and our Digital Footprint attack surface inventory tool. Also, many managed security service providers are offering a variety of digital threat management within their service portfolio. This would allow small companies to extend their digital defense capabilities.

Platform Solutions for Third-Party Cyber Risk Control

Risk management solutions for third-parties are emerging as a critical component of enterprise cyber security

Jonathan Dambrot, CEO of Prevalent

Few would argue that third-party security risk has emerged as one of the primary concerns for the modern CISO. Potentially exploitable weaknesses or incidents from vendors, suppliers, outsource teams, consultants, and other providers to the enterprise, have become as likely root cause culprits for major breaches as any other component in an organization. Traditional controls such as checklists and contract language have not been effective, and clearly, more advanced tools with better automation are required. Jonathan Dambrot, CEO of Prevalent, spent some time with us sharing his perspective on third-party risks and the best way to improve visibility and metrics-driven management into these potential weaknesses.

EA: Is third-party risk management primarily an issue for larger businesses?
JD: Any company that deals with suppliers, partners, and even business customers is dealing with third-parties. Even your accountant, lawyer, or auditor could be viewed in this manner. Our recent acquisition of Datum Security was intended to help us better address this issue in the SMB sector. So, the size of business you operate clearly does not dictate whether you deal with third-parties, but that size certainly has implications on the number and scope of external entities with involvement in your day-to-day activity. Larger companies can see partners and supplies in the thousands or even tens of thousands. The security implications of this type of arrangement are obvious.

EA: What are some of the more common third-party security risks that your team has seen?
JD: The biggest issue involves custody of sensitive data. That's gotten the most press in our industry, because it's been at the root of so many incidents. You know the situation: Your third party gains access to customer records or other important information, and then through sloppy handling or ineffective controls, allows that data to become compromised. It's your reputation that suffers afterward.

EA: What is the essence of world-class risk management? Is it visibility?

JD: The process of gathering and analyzing third-party controls is extremely cumbersome. What we do at Prevalent involves simplifying and automating this task with your vendor. In fact, we do it for your entire vendor network, so that gaps and seams do not exist in compliance or security. A major innovation at Prevalent is our vertical vendor network concept implemented as our Synapse Exchange™, which maximizes the collaboration between businesses and their third parties. A company can request information about a vendor, and if it is part of the network, then that information can be provided quickly, saving organizations a lot of time and resources to collect that data from the vendor.

EA: Do most of your customers generate third-party risk reports for executive review?
JD: Our ability to support the generation of third-party reports for executives is a major differentiator in the Prevalent Synapse platform. Accurate reporting is a powerful method for ensuring that all levels of executives in a company are aware of the risks from vendors. They get to see analyzed and synthesized output from assessments, questionnaires, risk scores, findings, and scans.

EA: What is the future, in your opinion, of third party risk? Do you think CISOs will get this under control?
JD: Risk from vendors has been such a serious security and compliance problem for so long that it is inevitable that CISOs will bring this under control. Our Synapse solution is designed specifically to help them accomplish that goal in a cost-effective and expeditious manner.

Managing Identity for Cloud, Mobile, and Premise

IAM solutions must adapt to the evolving needs of the modern enterprise using cloud, virtual, and mobile

Andre Durand, CEO of Ping Identity

Most enterprise security professionals would agree that identity and access management remains one of the most strategic aspects of their protection program. With integrations reaching across all aspects of an enterprise, IAM systems carry the burden of supporting identity provisioning, authentication, identity management, and enforce highly complex authorization and access requirements. Add cloud migration and mobile device management support into the mix, and the increasing complexity can be significant. Andre Durand, CEO of Ping Identity met with us recently to share his unique views and insights into the IAM marketplace, and how solutions are best offered for enterprise customers modernizing their infrastructure.

EA: Andre, do you see IAM as gaining in strategic importance in most enterprise environments?
AD: With the progression to hybrid cloud in the modern enterprise, the only way to create seamless control is to get the identity and access management approach right. The Ping identity platform is designed specifically to address IAM at scale across cloud, mobile, and legacy on-premises, becoming the primary control plane. Most IT and security teams have begun to realize this shift.

EA: What are the biggest challenges for IAM teams as companies migrate to mobility and cloud?
AD: I would say the biggest challenges involve maintaining the critical capabilities every enterprise requires to ensure seamless, secure access to applications including single sign-on, multi-factor authentication, access security, directory, and data governance. These are vital capabilities for any enterprise. And increasingly they must work seamlessly across a hybrid IT environment that span on-premises applications and cloud.

EA: How do you see identity federation evolving in the coming years? Will the larger cloud providers federate to everyone else?
AD: Yes, and great progress has been made already. With application portfolios spread across hybrid IT environments comes more complexity. Enterprise teams must work with modern platforms that support the latest standards to meet federation requirements between all participants including SaaS, on-premise and cloud environments. Users demand SSO at scale and without standards-based federation, that is not possible.

EA: What will be the role of mobile devices in future enterprise IAM?
AD: Mobile devices will become, and you could make the case that they already have become, fundamental components to the overall IAM solution. Every person and business is now dependent on mobility, so the integration of these tools, services, and infrastructure into IAM has become a key aspect of modern enterprise security.

EA: Tell us about some of the new capabilities and streamlined services being developed by your team.
AD: The Ping Identity Platform is continually striving to deliver a highly secure and seamless experience for end users and administrators. We see large enterprises adopting new digital business initiatives which also expose new risks and new gaps created in an organization's security, compliance, and user experience. We are extending our platform with new customer IAM capabilities, advanced multi-factor authentication features, data governance and regulatory compliance capabilities, and we are continuing to make it seamless to migrate from legacy IAM systems to the Ping Identity Platform. We are focused on helping enterprises close those gaps before they are exposed to risks while also giving their users a fantastic user experience.

Leading the Charge in the New Domain of Agile Risk Analytics

Simplifying and automating how enterprise teams manage cyber security risk through advanced IT inventory and security vulnerability management.

Nik Whitfield, CEO of Panaseer

Most CISO teams would list measuring, communicating, and tracking the mitigation of compound IT risk as their biggest challenge today. Why? Because while they have lots of data from individual security technologies and IT systems, they struggle to join these up and get a continuous view of what their data tells them about overall status and their next best decision to reduce exposure to compromise. After seeing that even mature and well-funded teams in the largest banks were struggling with this – and often dealing with a patchwork of business intelligence, operational alerting, and Big data tools to measure and report on risk, Nik Whitfield founded Panaseer. We caught up with him recently to ask about analytics, risk, and his views on the industry.

EA: Nik, why should CISOs care about so-called 'agile risk analytics'?
NW: Ultimately, this is about being able to show strong control over your exposure to compromise and impact, and your decisions about security priorities. As the CFO at one of our clients put it before they started working with us: "I need to be able to point at evidence based on good, current data and tell shareholders. This is what justifies my level of comfort about our risk posture and governance. And with cybersecurity, I can't do that." That quote isn't about solving all the problems we face as security teams; it's about being able to justify where you are, and where you're going. To do this, CISO teams need to be able to get many different views into their data, and spin on a dime when they get asked a new and often harder question by a regulator, an executive, or a customer. Today, firms lack that agility for several reasons; they don't have a dynamically updating inventory of their hardware, software, and people assets. That is, the operational underpinnings of critical business services and revenue generation. As a result, they can't look at alerts and information about the coverage and consistency of security control performance to assess

completeness of the picture they're looking at. And finally, they can't run that picture continuously to see the trends and systemic issues that tell them what needs to change across geographies, teams running IT platforms, and business units.

EA: So, this sounds like a real-time dashboard?
NW: So, yes – the reporting of meaningful, timely information in a visual interface is one aspect of it. But fundamentally, this is about what underpins that dashboard. It's about your security teams' ability to take a data set they've never used before, however large it is and fast it moves, correlate it or use it to enrich a data model, then analyze it to answer a risk question. And to be able to do that in days or hours not weeks or months. What this isn't about is automating PDFs into a dashboard. The risk problem in security is not slide deck automation problem. The problem is what's *in* the slide deck, and the fact it's generally 100 pages of metrics. The problem is that security is forced with its current toolset to abstract away the analysis and data that execs are presented with – things that it's critical for execs to have visibility into if they're to trust in the insight they're shown. The problem is getting a view into data that takes the multiple versions of the truth that security and IT teams are all working with across the data puddles they own, and pulling these into a single version of the truth that can be used to understand 'our next best action to get the biggest uplift in protecting our business'. And then being able to track that activity to understand how much risk is being closed out, how different teams are performing across risk treatment plans, and how many exceptions and exemptions you have across all your assets.

EA: Is that a big problem for a start-up to be solving all at once?
NW: Yes, and so while the technology platform we've built is designed to solve this at enterprise scale, that is, to be market leading – the problems it's solving today are 100% customer driven. And the number one problem, unsurprisingly, starts with 'device inventory'. It's not new, it's not sexy. But it's the foundation on which everything else is built – the accurate cyber hygiene metrics; the 'risk hunting' across workflow and security control data and logs; the threat detection analytics. While firms generally have at least one CMDB (or three), there's always varying degrees of accuracy. So, we take data sources from security controls and IT systems and mash that together to generate a dynamic picture of devices on the network. You'd be surprised what you can do just taking vulnerability scan, AV, SCCM, and active directory data – then comparing it with the CMDB. First you get a view of how complete your view of your environment is. And then you can compare vulnerability scan coverage to anti-virus coverage. Where are the gaps? Does it reflect your expectations of your security and IT process baseline that you thought was in place? And if not, what teams do you need to talk to tomorrow? It sounds basic, but as a CISO client of ours says, it's not the super-mega-advanced stuff that causes most breaches. It's the simple stuff you look back on and think 'If only I'd known, I could have stopped that easily!' Oh, and by the way, now take that problem and imagine it across all your 300 or 3000 suppliers!

EA: Why did you decide to found a company to solve this problem?
NW: At BAE Systems Detica, the founding team and I were among the first people to commercialize data science for threat detection problem sets. We were dealing with scenarios where a CISO would have been given an IP address by an intelligence agency, and

needed to do proper needle in haystack stuff with big, hairy number crunching on data sets like web proxy. But what we consistently came across in those teams is they'd tell us 'We spend 60-70% of budget on detect and respond; we're seeing diminishing returns; we need better security, not more alert triage'. When we founded Panaseer we made a very conscious decision not to do 'threaty stuff'. Not because that wasn't valuable, but because we knew that firms were looking at the NIST-aligned areas of Identify and Protect, and saying 'How do we use our data to get upstream of the problem and optimize our security so detect and respond can be pointed with laser focus at things where we can tolerate a high false positive rate because we have a much smaller scope, rather than detecting everywhere, but only seeing white noise?' I mean, when getting ten thousand false positives a day from a SIEM solution is regarded as good, you know there's a need to fundamentally change how we're doing things. So, we looked at the problem differently and very quickly we knew we had something because CISOs jumped on our technology because it enabled them to shift their approach.

EA: So, what's next for Panaseer?
NW: We've been working with global financials for the last 2 years in both London and New York City. We opened our New York City office last year and now we're expanding in the US market. What that means in practice is building integrations with partners, particularly workflow tools to support operational automation of 'action on risk'. We're growing our engineering team too, as we're seeing huge market demand for solving this problem. So, it's all about making sure that we can make Panaseer available to more clients and keep helping the security industry develop in how we go about protecting the companies and the infrastructure we all depend on. The security has had to evolve at the speed of the internet, and so we're innovating to help security teams manage the risk that comes with at the speed of data.

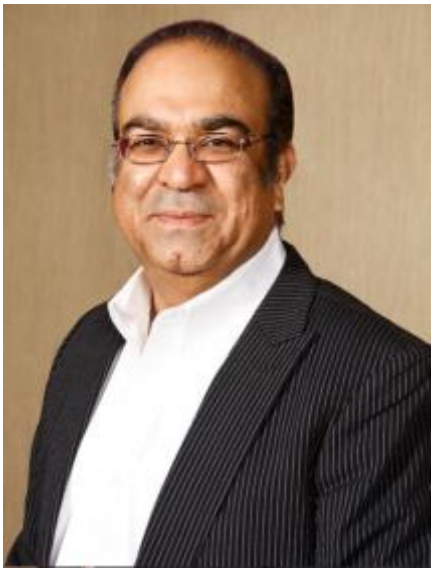

Providing Rapid Cyber Analytics for Data in Motion

Modern enterprise and critical infrastructure protection demand advanced data capture and processing tools

Dr. Parag Pruthi, CEO of NIKSUN

Few companies understand the requirements for high-speed packet capture and analytics-based processing to detect cyber indicators like NIKSUN. The company has been at the forefront in this area for many, many years, and its founder and CEO, Dr. Parag Pruthi, has been improving technology platform solutions in this area for decades. We sat down with Parag recently to ask him to share his thoughts on how platform design is evolving, as well as how the underlying behavioral analytic algorithms are improving to the point where they can dependably identify real cyber attack indicators in enough amounts of collected packet data.

EA: What types of packet capture features are your customers requesting?
PP: From the time I started working in cyber security and founded NIKSUN, our message has been consistent, loud, and clear – Murphy's Law paraphrased as "the packets you did not have were the packets you needed" is well and alive. The business implications of this empirical rule are profound. That is, the very tools you don't have are just the tools you needed to do the job right. As a result, the number one feature our customers request is this: "Don't lose any packets because I don't know when I will need the one that were lost." However, for NIKSUN, this is a given – zero packet loss at 1Gbps, 100Gbps, 1000Gbps, or whatever rate is desired. The second feature concerns help with the question: "How do I find the needle in the haystack?" To satisfy this request, at NIKSUN, we first index everything, from the packets and the data in those packets to the sessions, and the data in those sessions, to the applications and the data in those applications. Next, we provide a single portal for analyzing all this data from anywhere and at any time. That is, whether the data is in the cloud, in different data centers, in a virtual environment, or scattered across an enterprise over multiple asymmetric routes, it all just needs to be accessible to an analyst in the same way from one place, irrespective of the analyst's physical location or device. Last, a third and an increasingly important feature requested by our customers is

support for easy and fast visualization. That is, they are looking for help with the question: "How do I know what to look for?" Now, satisfying these, and many other requested features, and making it all work while the application landscape underneath you is constantly changing is no easy feat. Having done this well is the only reason why NIKSUN is the solution-of-choice for not only the U.S. Department of Defense, but also any enterprise for which actual results matter more than marketing hype or personal connections.

EA: Are behavioral analytic algorithms efficient enough to perform sufficient processing for real-time networks?
PP: Well, it depends! Some analytic algorithms are efficient and others are not. The distinction lies in understanding their practical use in cyber security. For example, we can perform principal component analysis in real time and build a language to form expressions of those components. Under certain conditions, this method works rather well and can capture known anomalies where signature detection would fail. Also, deep learning methods, such as deep or recurrent neural networks, are in vogue today, and I am often asked if they can be used to find zero days and stop all cyber attacks. On the one hand, machine learning (ML) works well in some domains where classification is rather straightforward and ample training data is available to converge the algorithms at the minima and not get stuck at false valleys. At the same time, without significant commonality in the various attack vectors and the lack of sufficient training data, all behavior analytic algorithms (ML-based or not) need to narrow the analysis using a variety of depth or breadth first search algorithms. The resultant state space can become so complex that it is not possible to do so within practical budgets. Thus, while there exist algorithms that are specifically devised to detect certain anomalous conditions and are amenable to real-time analysis, many of them are unfortunately not yet computable in real-time. At NIKSUN, we develop both real-time and non-real-time expert systems which encompass various algorithmic analytic techniques.

EA: What is the accuracy of typical analytic algorithms in detecting threats on high-speed networks? Is the false positive rate low?
PP: Despite the numerous success stories of purposefully-designed ML-powered artificial intelligence systems – for example, all of Facebook's translations are now completely powered by an unsupervised deep learning system – their applications to cyber security have been less than stellar. One of the main reasons is that in cyber security, a key challenge is the detection of "unknowns" (i.e., threats never seen before) in close-to real-time despite often very weak signals. For various reasons, this is a task at which unsupervised ML does not excel. Outward signs of this mismatch between what the cyber security domain demands and what unsupervised ML techniques are good at are unacceptably high false positive rates that severely limit the use of ML-based AI systems in practice and unreasonably large mean dwell times that all but guarantee that the attackers have the luxury to take their time to achieve their various objectives. The basic problem with using such systems is that once the analysts lose faith in them due to the high number of false positives, they tend to ignore all the generated alerts and fall back on what works for them – performing everything manually. They react similarly when faced with detection times for breaches that are measured in days and months. Thus, the net effect of using such systems can be self-defeating when applied to the domain of cyber security without proper

restraints. However, for carefully designed systems that are applied with the proper restraints, the cyber security domain provides enormous opportunities. For example, when using the fundamental approach (i.e., "the NIKSUN way") of collecting and indexing all the data and combining it with both algorithmic techniques and computer-assisted but human-navigated analysis, the results can be remarkable. By experiencing efficiency gains far exceeding 500% over traditional methods, many of our clients can do significantly more work with fewer people; or in other words, their analysts can focus exclusively on getting the upper hand over the bad actors.

EA: Do you see changes in the mix of hardware and software required to provide advanced analytics at line speed?
PP: Well, it depends on the line speed. At low speeds, software-only solutions will suffice. But at very high rates, software alone on general purpose hardware is inefficient and impractical. Somewhere in between, a mix can be leveraged. My point is that to be able to defend oneself against attacks such as the recent zero days like WannaCry and Petya as well as a host of other more complex cyber-attacks, one must consider all possible known or unknown attack vectors. As a result, one has no choice but to deal with the problem head-on. Basically, advanced analytics at line speed poses three big challenges. We already talked about the technical problem of performing lossless packet capture and simultaneously generating associated meta-data at high speed (e.g., 1-100 Gbps and beyond). Next, today's cyber security is all about close-to-real-time detection and mitigation of nefarious activities, with the added feature of being able to perform retrospective network forensics when needed or required. This overarching desire for real-time solutions upends traditional analytics and requires the collected data to be treated as streaming data where any analytics is based on a one-time exposure to the data (i.e., at the time of data capture). Essentially, batch processing techniques need to be reinvented to work in real-time. A further challenge is posed by the distributed nature of a typical modern enterprise network. In fact, the ability to collect high-velocity and high-volume streaming data at different locations in such environments mandates a fundamental shift in data analytics. The traditional view that "the data has to be moved to where the analytics/processing is done" is replaced by the new insight that it is the analytics that must be distributed (i.e., the analytics/processing has to be brought to where the data resides). Even though these challenges have been known for the last 20 or so years, some people still wrongly think that a simplistic mix of solutions can address them. For example, one popular solution is to have a device classify all the data, such as a firewall, and another one to simply collect packets. The problem with this approach is that doing the real-time analytics or post-event analysis without the appropriate metadata is by and large useless. By the time one is done fetching the packets for specific flows and then reassembling them for analysis, many other events will have queued up and there is no digging out of this hole. At NIKSUN, we have studied this problem very carefully and ended up designing a solution that we optimized in both space and energy, in hardware as well as in software. NIKSUN's Supreme Eagle architecture, with its built-in support for cluster and grid computing, provides exactly the type of system-level support that this paradigm shift in advanced analytics requires. As an all-in-one platform, it offers the basic functionalities for real-time analysis of the type of "hyper data" that it collects. In fact, it is ideally suited for implementing distributed streaming data algorithms that are at the core of any advanced

analytics in support of real-time cyber security solutions. We have advanced this mix of hardware/software analysis so far that we are now exploiting the full power of this type of advanced analytics to harness unprecedented opportunities for both real-time cyber security solutions as well as "back-in-time" analysis. Moreover, by supporting this type of advanced analytics on our suite of virtual solutions, we can offer customers "network monitoring as a service" and enable them to reap the benefits of network function virtualization (NFV) by letting them decide where to perform ultra-high performance packet capture and analytics, when, and for how long. NIKSUN's virtualized software takes full advantage of dedicated hardware and provides scaling in multiple dimensions.

EA: How important is domain knowledge to detect network attacks for applications such as industrial control or IoT?
PP: If the past is any indication of what the future in cyber security has in store, we would be foolish to envision that we will be able to completely replace domain experts or eliminate humans from the loop by leaving it all up to AI to do the job for us. Whether we are concerned with protecting the various systems that control the myriad of different industrial organizations and critical infrastructure networks we rely on in our daily lives or worry about nefarious activities that potentially involve millions of vulnerable IoT devices and can presumably cause havoc at local or global scales, domain knowledge will remain the go-to solution so long as the software for the control is written by humans. Just as domain knowledge is paramount for finding bugs in this software, recognizing how they can become vulnerabilities when used for nefarious activities, and ultimately exploiting them for specific attacks, it is also essential for reverse-engineering a given (unknown) bug from an observed attack. While AI in its current form is ill-suited for both these tasks, domain experts excel in them. At the same time, once the basic mechanisms underlying such "unknowns" have been elicited and are understood and known, the job of detecting future occurrences of the same type of attack in real-time can be left to AI after the successful implementation of suitable real-time analytic algorithms that mimic the steps used by the domain expert to get to know these unknowns. It is in this sense that existing and emerging AI approaches can be fully expected to play a critical role in securing our future networks against cyber attacks. By automating all the tasks that are amenable to automation, we reap the benefits of AI systems by putting ML techniques to work on problems where they reign supreme – detecting "known bad" activities with high confidence and preventing "known good" activities from triggering false alarms. At the same time, this use of AI also frees up the domain experts to focus on work where they excel at – getting to know the unknowns in a gradually diminishing portion of suspicious traffic. It is in this sense that I believe that the holy grail of cyber security – that is, the real-time detection and mitigation of nefarious activities – will for the foreseeable future require human involvement in the form of cyber security experts and their invaluable domain knowledge.

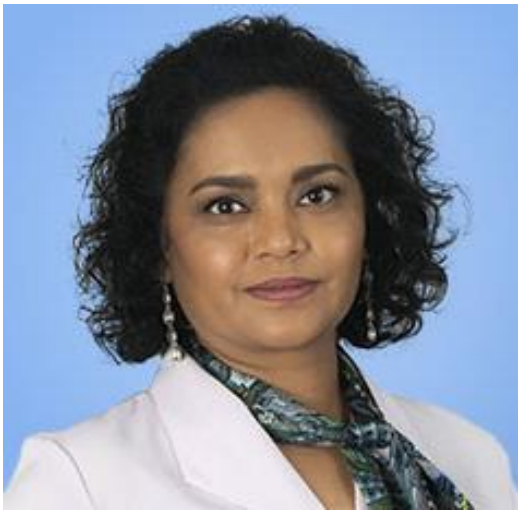

Using Isolation to Reduce Malware and Attack Risk

Cyber attacks such as phishing can be isolated to virtualized cloud-based platforms for safe security filtering

Poornima DeBolle, Chief Product Officer of Menlo Security

The endpoint has become increasingly vulnerable to a wide range of different content-based attacks from websites. Pointing your computer at a public website such as CNN.com, for example, causes a plethora of different scripts and executables to rush down onto your computer via the browser, and expose you to serious known and unknown risks. Cyber experts have recently identified an effective man-in-the-middle solution where such sites are processed by isolation platforms in the cloud to filter possible attacks. The result is then remotely rendered to the user, who has no degradation in experience, but no longer has the risk exposure. Poornima DeBolle, Chief Product Officer of Menlo Security, has helped to pioneer this approach, and she sat down with us to provide her perspectives on the prospects for such technology in the enterprise.

EA: Tell us how isolation works for endpoint security?
PD: Websites and email remain the two leading vectors for malware to reach the endpoint, and the threat is now so prevalent that IT organizations and individuals fear clicking on any active web content or email links. The cloud-based Menlo Security Isolation Platform (MSIP) eliminates the possibility of malware reaching user devices via compromised or malicious web sites, email or documents. This is not detection or classification, rather the user's web session and all active content such as Flash, whether good or bad, is fully executed and contained in the Isolation Platform. Only safe, malware-free rendering information is delivered to the user's endpoint. No active content – including any potential malware – leaves the platform. In the Menlo Security isolation model, malware has no path to reach an endpoint, and legitimate content needn't be blocked in the interest of security. Administrators can open more of the Internet to their users while simultaneously eliminating the risk of attacks.

EA: How does the rendering work? Do users see differences?
PD: Menlo Security's patented Adaptive Clientless Rendering (ACR) technology provides the connection from the user's session running in the MSIP to the user's native browser.

For each type of web content the ACR engine selects the optimal encoding and transport mechanism for delivery to the user's browser. For example, dangerous content such as Flash is executed in the MSIP and then delivered as a hi-fidelity, interactive experience in the user's browser. In all cases, the user's browser receives non-executable, malware-free content that renders naturally. ACR technology requires no endpoint software or plug-ins and delivers a completely native user experience essentially indistinguishable from direct interaction with a web site.

EA: How do you protect your own platform in the cloud?
PD: Menlo Security recognizes that malware protection not only depends on the efficacy of our products, but also on the security of our infrastructure. The MSIP and management console are hosted in Amazon Web Services (AWS), a secure environment that is continuously audited, with certifications from accreditation bodies across the globe. In addition to AWS native defenses, Menlo Security has gone to great lengths to protect customers. This includes extensive cloud infrastructure security, and regular third-party platform security audits. The MSIP creates an isolated browser instance for each tab in a user's browser session. That browser is created inside a hardened Linux container which includes only the resources needed to run a browser session. Within each container, the browser process is subject to a mandatory access control policy governing resources available to the browser. Any attempts to exceed container limits is blocked and alerted. Each container and browser is then destroyed at the end of every web session.

EA: What's been your experience working with enterprise teams regarding isolation? Have IT teams learned to appreciate the risk reduction?
PD: CISOs are generally aware of the risks coming from the web and email. Some of our customers started on the path of isolating only web sites that were tagged as "uncategorized" by their web gateways. Once they saw the risk reduction from just eliminating that threat vector, many quickly moved to isolate all web traffic, as they recognize the potential to eliminate entire categories of risk. We have worked with several enterprise customers who have pegged that risk reduction at over 85%.

Cyber Situational Awareness via Visibility and Analytics

Enterprise security benefits from platform automation with advanced threat analytics

Sanjay Raja, CMO of Lumeta

Experienced enterprise security professionals understand the value of situational awareness to address modern cyber threats. Enterprise security risk has shifted from a compliance-driven response obligation to a risk-driven proactive challenge. To support such continual protection, platforms are required that combine the best elements of visibility support with advanced analytics to offer an accurate picture of on-going vulnerabilities and potential solutions. Sanjay Raja, CMO of Lumeta, spends considerable time thinking about this topic and he sat down with us to share his unique insights.

EA: What is the benefit of cyber situational awareness for enterprise security teams?
SR: Anyone with responsibility to protect an enterprise knows how important it is to have an accurate understanding in real-time of all your assets and infrastructure, as well as changes to that infrastructure that are indicators of what can lead to malicious activity. We start out identifying all your unknown, unmanaged, rogue, or shadow IT infrastructure. This includes networks and endpoints, both physical and virtual. Most successful breaches and ransomware attacks can clearly be traced back to a lack of immediate visibility or awareness of existing network, networked, and cloud infrastructure leading to the inability to account for and protect systems proactively, as well as missing malicious network activity due to the limitations of existing security stacks. These elements create the underlying motivation for our platform offerings at Lumeta, where we strive to provide accurate analytics-driven visibility of the security posture of the entire network to security teams for optimal mitigation and response

EA: How does your platform collect data to create visibility across an enterprise?
SR: Our flagship product, Lumeta Spectre, combine a set of patented active probing and passive listening techniques at the network layer extending all the way to the endpoint and into the cloud. Customers working with our powerful solutions are reporting on average,

40% reduction in so-called 'blind spots' in their infrastructure. As you know, it is these blind spots that lead to the most serious intrusions. In addition, Spectre provides a real-time understanding of changes in the network, but we don't stop there. We pull in threat intelligence that is also applied to our network flow modeling to develop what we call, Threatflows. These flows are indicators of malicious behaviors on the network, whether flagging compromised systems, identifying leak paths to external malware hosts, or identifying encrypted communications like TOR, that are often not authorized for use.

EA: Can you share how you approach analytics?
SR: We take a unique approach that allows us to provide 100% coverage across the network versus existing methods that leverage packet captures, logs, and netflow, all of which provide an incomplete picture due to the limits of those technologies. At Lumeta, we focus on the underlying network infrastructure that forms the basis of all communications in the environment and is central to discovering attacker activity. Since we look at primarily the network control plane traffic, we can discover recursively, and collect and analyze every network, networked device, and/or endpoint. Our analytics differentiates between the relevant protocols such as OSPF, BGP, ARP, DHCP, DNS, and ICMP. Protocol-specific information is rapidly correlated with discovered contextual data on the network, endpoints, and across the hybrid cloud to detect relevant changes to the infrastructure. The types of changes our analytics identify include new bots, new C&C points, newly accessible Tor exit nodes, unusual port usage, and many other focus areas. The emphasis is on speed, accuracy, and relevance to cyber security concerns. As we apply threat intelligence to relevant metadata, our analytics can provide areas of vulnerability to attack, but also indicators of compromise and potential breach activity.

EA: Does your solution support network segmentation?
SR: Absolutely. As we have a complete understanding of the network and changes in real-time, we can search for so-called 'leak paths' between presumably isolated segments or even leaks to the Internet, including from the cloud. This knowledge can be essential to ensuring proper segmentation security and compliance, and determining if violations exist such as undesired lateral movement or unauthorized communications, especially to the outside. When we detect such threats, we can provide this data in real time to the SIEM or other collection device in the enterprise. Our ability to provide this type of information can also help network and security teams accelerate their "unflattening" of networks and optimize their segmentation to ensure it is configured as expected, but also flag violations.

EA: Which business sectors or industries will benefit most from such visibility capability?
SR: Obviously, critical infrastructure sectors have the most intense obligation to support advanced real-time security controls. So, we've seen great focus from these larger companies, agencies, and organizations. But more recently, we've seen middle and even smaller sized businesses paying closer attention to real-time visibility on the network, endpoints, and cloud. This includes all sectors such as financial services, retail, technology, telecommunications, services, and on and on. We've discovered that there is really no size or type of business or government agency that will not benefit from our capability. One growing segment has been IoT, but focused on areas like health care, manufacturing, utilities, retail, and critical infrastructure. Our solutions have proven to rise above the hype

around IoT security as our core platform is the perfect foundation for providing visibility, securing systems more effectively and detecting leaks or other attack activity, while the promises made by some vendors takes years to bridge the gap between the reality and hype.

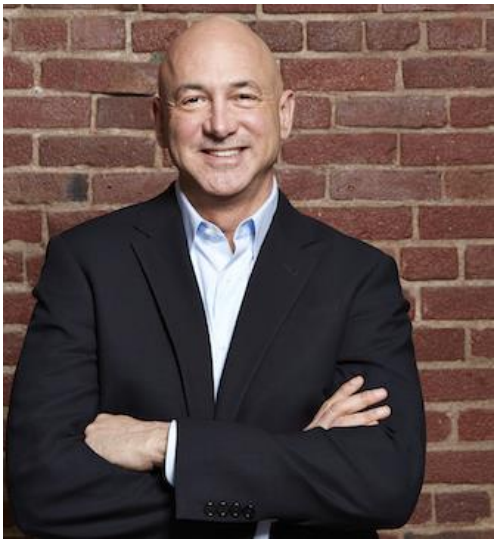

Securing Mobile Devices and Identity

Advanced protection of mobile devices and user identities can be achieved using a single mobile application

Jim Dolce, CEO of Lookout

Mobile devices are a part of every enterprise's critical infrastructure. Employees use them every day to work. These devices access significant amounts of sensitive data and act as a conduit, transporting that data off the device through email and applications. Threats, vulnerabilities, and other risks to data that affect PCs also apply to mobile endpoints, yet simply extending current PC security controls to mobile is ineffective. Security professionals must redefine their approach to risk management in the mobile world, and architect mobile-specific security. I had the pleasure to meet with my friend, Jim Dolce, CEO of Lookout, recently to discuss advances in mobile security and related trends in modern cyber risk.

EA: Jim, do you see both enterprise users and consumers recognizing the value of advanced mobile protections?
JD: Yes. A decade ago, when Lookout got its start, mobile threats were primarily commodity malware or SMS scams, issues that consumers needed protection against. We built a strong consumer business, which now protects more than 100 million devices and we continue to see interest from our carrier partners in expanding the security they can offer to their mobile subscribers. Today, our enterprise business is also thriving with considerable uptick happening in just the last twelve months. As targeted mobile attacks have become pervasive and relevant to the enterprise, as regulatory bodies begin to recognize the amount of information accessible via mobile devices, and as CISOs struggle to control corporate data when employees choose which mobile apps they download and use, we are seeing more enterprises realize that mobile is the forgotten endpoint in their security strategies.

EA: What sort of algorithms are used to detect threat to the mobile ecosystem?
JD: Within the Lookout Security Cloud, we utilize a range of algorithms to detect complex patterns that indicate known and novel threats, software vulnerabilities, and risky

behaviors and configurations. Our cloud-based machine intelligence conducts fuzzy comparisons to the app characteristics in our dataset allowing us to see connections, correlations, and relationships that simple signature, network, and behavioral analysis based systems routinely miss. What's most important here is that none of this technology would mean anything if our dataset wasn't large. In fact, with more than 100 million devices and forty million apps and counting in our network, we have the largest mobile dataset in existence. This data allows our machine learning algorithms to accurately identify anomalies and make connections in code that would otherwise go unseen.

EA: Do you see improvements in the security functionality supporting mobility? For example, do you see OEMs and ISPs doing a better job patching?
JD: Yes, efforts by Google, Samsung and others to accelerate patch cycles have made a considerable difference in addressing Android OS vulnerabilities. However, an area where we need to see broader ecosystem prioritization of security is app vulnerabilities. At Lookout, we repeatedly see issues where a developer has unknowingly included flawed code or integrated a risky SDK or a compiler into their app, thus exposing users to possible surveillance or data exfiltration. Some of these apps make it into the popular app stores and some are built by enterprise app development teams. We need all app developers to build with security in mind.

EA: What are some mobile hacking trends that your team is seeing?
JD: Enterprises take a lot of care to protect their laptops from APTs. But in the last year, we've found evidence that the same nation-states enterprises are spending millions to keep out of their laptops are widely executing targeted attacks on mobile. Much like we see in the desktop malware space, socially engineering users via unsolicited messages, in many cases SMS, is an effective approach for these attackers. And while targeted journalists and political dissidents have been in the headlines, there's growing evidence that enterprises are being targeted too.

EA: Have audit and regulatory authorities begun to recognize the need to include mobility in their compliance frameworks?
JD: I'd like to see the compliance and regulatory communities move much more quickly and forcefully in securing mobility. They've made improvements in the applicable control frameworks, but mobility is no longer an add-on to typical enterprise infrastructure, it is the backbone to most operations. Major regulations, such as SOX, GDPR, HIPAA, etc. must recognize the risk that is exposed when mobile is unprotected.

Rapid Network Threat Detection and Analysis

Advanced threats can be detected on networks via real-time controls and behavioral analytic algorithms

Brett Williams, CSO of IronNet Cybersecurity

Despite massive distribution and virtualization of enterprise computing, infrastructure operators continue to recognize the on-going and *increasing* requirement to factor scope, scale, and speed issues into cyber security architectures. To keep up with these functional demands for modern infrastructure protection, the cyber security team must take advantage of every possible technical and mathematical approach to detecting unknown threats and initiating rapid, automated response. Brett Williams, CSO of IronNet Cybersecurity spent some time with us recently to share his thoughts on how cyber security defenses can keep up with rapid and powerful offensive attacks.

EA: Brett, your team uses the term collective defense – what is meant by that?
BW: To begin, collective action on the offensive side already exists and is a key factor in the ability of cyber adversaries to continually outpace cyber defense. At the grass-roots level, vibrant black markets for vulnerabilities, cyber-attack kits, botnets, and other offensive capabilities are readily accessible to anyone with the right browser plug-ins to reach the anonymous forums that house these markets. At the high-end of cyber offense, proven capabilities are intentionally and unintentionally shared by governments with less sophisticated criminal and activist groups to leverage against new targets of interest. Social media sites serve as powerful amplifiers, spreading knowledge and source code for the development of enhanced variants of malware and offensive toolkits. Cyber defenders are overtaxed by the sheer volume of information they must deal with and suffer from a skills gap. No organization, no matter how large, has the resources to stand alone against nation-state threat actors, let alone the myriad of real and potential criminals, hacktivist or other threat actors targeting their organization. To stand alone is to be divided and conquered–to fall one by one. As a result, collective defense is an absolute necessity to cope with the present and future cyber threat environment. Companies need to band together in common defense to gain broader situational awareness of the threats targeting their specific sector, and to jointly mitigate threats targeting the collective group. For collective defense to

function properly, information sharing needs to occur at network speed and across a broad base of indicators, risk-models, and enrichment resources. With such a system in place, an attack on any organization in the collective can be immediately addressed by all.

EA: How important are advanced analytics to the detection of unknown threats?
BW: Cyber security has become a data aggregation and mining problem. Time to detection of threats is critical to any cyber defense and is increasingly difficult given the complexity and noise within an enterprise's network environment. The goal of leveraging advanced analytics for detection is to move beyond identifying moment-in-time events to modeling adversarial tools, tactics, and procedures (TTPs) used to orchestrate and manage attacks. Detecting offensive TTPs can help reduce overall cyber risk by shrinking the offensive playbooks available to the adversary, thereby reducing the potential and efficacy of available tactics. Producing adversarial models is probably the most difficult level of detection for cyber analytic solutions as it requires a non-trivial investment in time and resources to develop defensive models against the most likely threats to an organization. It requires access to high quality data scientists and experienced cyber defenders with expert knowledge of the tactics used by advanced threats. However, when done properly, the benefits to such an approach is substantial. It enables detection across an adversary's full range of tactics and not just at the late stages of the kill-chain. This improves the overall detection capabilities and it raises the bar for the offense by forcing them to have to design new targeted TTPs as opposed to simply retooling existing approaches with repackaged malware, scanners, and tools.

EA: Do you think there is much that commercial industry can learn from government teams regarding cyber defense?
BW: Absolutely. Defending against nation-state threats across vast computer networks in the federal government has resulted in advancements in defensive tactics, operational procedures, and detection techniques that are at the cutting-edge of cyber defense. Equally important, the highly-trained personnel that have experienced both the defensive and offensive side in government are invaluable resources when they decide to move into the commercial sector. If you look at how commercial organizations are investing in cyber security today, you will notice that many in the commercial sector agree with this premise. Many of the new CSO/CISOs and their staff have had previous experience in an operational background defending government networks. Much of the security best practices or security frameworks that drive cyber security spending are based on government frameworks or derivatives of frameworks such as NIST. Many of the new cyber security start-ups that are developing new cyber defense products or services are founded by former employees of the US military or intelligence agencies and backed by blue chip venture capital firms. What I will also mention is that government teams have also learned from commercial industry with regards to understanding their perspectives and their businesses within the context of cyber defense initiatives. This helps government address some of the gaps or potential gaps that exist today with regards to cyber defense on national level. More importantly, a close partnership between commercial enterprises and the government will help improve the cyber resiliency of the individual organizations as well as the nation.

EA: IronNet Cybersecurity has always focused on the engineering requirements to keep up with attacks on high capacity networks. What is the secret here? Is it hardware? Software? Perhaps a combination?
BW: It's really all the above—plus data scientists and security analyst that can provide feedback from an operational perspective. Our solutions have always focused on empowering our customers' security analysts to be more effective across the full spectrum of their work as opposed to creating point solutions for small subsets of the overall problem. Consequently, our efforts have been to deliver the full range of capabilities necessary to support a security team's strategic objectives. The hardest part of building an effective solution, is having the expertise to put it all together at the size and scale necessary to support the mission. While many solutions leverage similar technology platforms or foundational techniques, in practice it can be very difficult to deliver an effective analytical solution with the breadth, scale and depth to defend against the full range of threats targeting the organization. Perhaps the easiest way to understand this is to think about the differences between a great chef and an average cook. Both start with similar tools—a knife, a pan, a stove and some ingredients. The difference is that the experience and choices made by a great chef can result in a dish that can be an order of magnitude better in taste that can scale up or down as needed to meet the size of the dinner party. It is the same way when building cyber analytic solutions — an experienced team with deep domain experience can build a cyber analytic solution that is an order of magnitude better in terms of detection capability, scalability, and user experience.

EA: What sort of trends is your team seeing from the most capable offensive actors?
BW: We are increasingly seeing the use of cyber attacks by state actors to project national power or national goals against countries, companies or other organizations. At IronNet, we are working closely with commercial organizations in critical infrastructure sectors to help defend their networks against all sorts of cyber threats. One recent case was the use of destructive malware against a subsidiary of Fortune-100 company located in a geographically sensitive region of the world. In this case, the malware was hidden in software distributed by the local government to be used by corporations for tax calculation purposes. The malware was particularly virulent and spread quickly throughout the organization, wiping out many of the computers, and resulted in material impact to the organization. We suspect that this was led by a nation-state actor trying to disrupt commercial activity to advance their national goals.

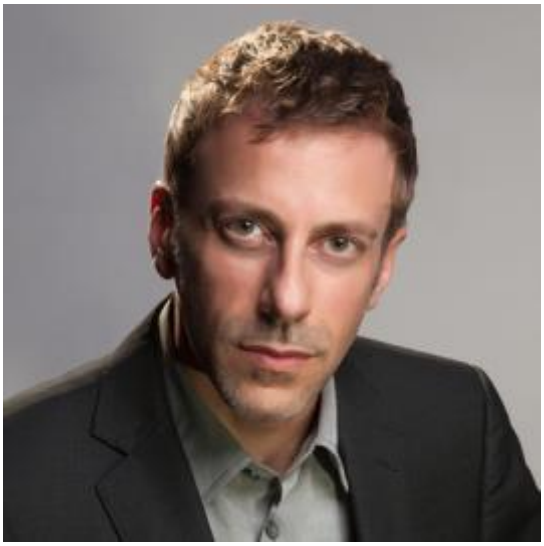

Extending Cyber Forensic Advances to the Endpoint

Security detection and response solutions for investigators can be effectively integrated into modern endpoint controls

Anthony Di Bello, Senior Director of Product, Guidance Software

Forensics has always been primarily about establishing visibility into the hidden or obscured. Uncovering fingerprints, analyzing striation patterns on bullet casings, and the use of ultraviolet light to reveal blood splatter are examples of this in the physical world. As it relates to digital forensics, investigators will tell you that the challenge in any case is unraveling the complexities of software, devices, networks, and systems to *make visible* the evidence, data, or information of interest. EnCase from Guidance Software, has long been at the forefront of the cyber forensic profession. Skilled investigators can now make pretty much anything visible that they focus on – and this is directly attributable to the advanced tools at their disposal. This level of visibility and investigative support presents a unique opportunity in detection and response when applied to enterprise endpoints. Cyber security experts have long known that the primary goal of endpoint protection involves visibility-based control. So, the extension of forensic advances in this direction is natural and welcome. Anthony Di Bello, Senior Director of Product at Guidance Software, was kind enough to spend time sharing his perspectives in both investigation support in cyber

EA: Let's start with an overview of the cyber investigative community. What sort of advances are you seeing in cyber forensics?
AD: The primary shift that we see involves the introduction of new devices and applications as sources of data and evidence. Smartphones, drones, and IoT devices are just a few examples of new devices types that may store information crucial to a case. As such, investigators need tools to quickly gather information from a wide variety of devices and operating systems. Building on nearly 20 years of experience at Guidance Software supporting the forensic investigator, we've expanded our product set to support the widest range of modern devices and the applications that run on them.

EA: Tell us how your forensic solutions have been extended to endpoint products.

AD: Our EnCase platform's powerful support for the forensic investigator gave us a natural ramp to an endpoint solution that could scale to the largest of enterprises. Most endpoint security tools, for example, constantly collect and send reams of data from distributed endpoints to security analysis platforms, much of which is irrelevant. This works fine for small numbers of endpoints, but does not scale to larger deployments. Our knowledge of surgical efficiency in detecting indicators in large volumes of data provides a powerful underlying capability to scale detection and response capabilities effectively.

EA: What are some of the techniques used in your products to make cyber security indicators visible?
AD: In the past, it was simple correlation. Now, we utilize advanced heuristic methods based on data analytics and other techniques. Quickly focusing in on cyber security indicators no matter how well hidden is the essence of the cyber forensic investigative task, and we have many years of experience assisting investigators and responders in that area.

EA: What sort of new hacking techniques do you see being used now by advanced malicious actors? Is it correct to assume that bad actors are getting much better?
AD: Bad actors are more organized. We see such organization expressed in trends moving from simple data breaches and commodity malware, to more complex and destructive malware. Ransomware, IoT-based attacks, as well as sophisticated zero-day attacks on the enterprise are all good examples. This shift to more serious and consequential cyber threats underscores the importance of having a world-class endpoint detection and response solution deployed across the entire set of endpoints.

EA: What trends do you see in the coming years in cyber security and forensic investigations?
AD: For cybersecurity, we will continue to see advanced attacks capable of penetrating perimeter security defenses as well as not-so-advanced attacks that take advantage of unpublished or unknown software vulnerabilities. For that reason, we agree with industry analysts who predict massive growth in the EDR market. Especially at the enterprise level, security teams will need the ability to quickly identify, validate, triage, and remediate threats to the network. For forensic investigators, they will continue to face challenges dealing with an ever-expanding universe of devices and apps and motivated adversaries. Powerful tools with deep visibility, robust support, and powerful remediation capabilities will be a must-have.

Enterprise Risk Reduction Using Security Fabric

Integrated cyber security fabric ensures advanced protection for enterprise with no interface seams

Phil Quade, CISO of Fortinet

The challenge of enterprise security is certainly not for lack of vendor tools. In the past two decades, a plethora of different solution options has emerged in every conceivable aspect of software, IT, and network infrastructure. What has been missing has been the ability to coordinate and orchestrate these tools, with the goal of maximizing synergy and minimizing seams between components. This is best done, in most environments, with a single pane of glass management console. Fortinet has been a leader in providing a suite of advanced solutions that enable enterprise customers to create that seamless fabric across their network infrastructure. The concept of a woven fabric to protect assets should resonate with any CISO who has struggled with varying platform interfaces and non-interoperability between vendors. We had the opportunity to sit down with Phil Quade, CISO of Fortinet, to discuss security architecture architectural issues, and how the Fortinet concept of security fabric can be such an advantage to enterprise teams.

EA: Phil, please start by telling us about your wonderful background in the Federal Government and how you are leveraging this experience to commercial solutions.

PQ: Threats to national and commercial security come in all forms, and my specialty is understanding how to strategize, plan, operate, and communicate prudent cyber security solutions. In my 30-year career at the National Security Agency (NSA), I've had the opportunity to work across the defense, intelligence, and attack aspects of cyber in domestic, foreign, government, commercial and critical infrastructure sectors. I worked most recently for the Director of the NSA, representing the agency at the White House and coordinating cyber efforts at the NSA. I previously served as the Chief Operating Officer of the NSA Directorate that focused on securing America's most sensitive and classified government systems. Before that, I served as the Head of the Information Operations Technology Center's Advanced Technology Group, as a professional staffer to the U.S. Senate, and at the Office of the Director for National Intelligence. Along the way, I've had some great assignments as a computer and network security evaluator, cryptanalyst, and

export policy specialist. At Fortinet, I enjoy applying my experience of managing diverse and complex cyber strategies with a variety of public and private partners to ensure that both Fortinet and its global customers have the most effective, broad security postures.

EA: In your opinion, what are the big challenges today in cyber security?
PQ: We are in a computing revolution, one that has its roots in the rise of mobility, the cloud, and now the Internet of Things. The adoption of a digital business model is requiring networks to evolve more rapidly, and applications, data, and services to flow faster across an increasingly diverse landscape of users, domains, and devices. As a result, today's networks and related security are also increasingly borderless. IoT and cloud applications, data, services, and infrastructure now require organizations to worry about an attack surface that may not even be visible to IT. We also face a huge volume of cyber threats, along with highly sophisticated targeted attacks, made possible by the commercialization of a whole ecosystem of cybercrime services and supply chain resources and services. Whether you have the right protection immediately responding to threats throughout your network can determine if your business runs smoothly or is the victim of a security breach. The cyber security risks individuals and organizations face today are real, and they come from vulnerabilities, threats, and lack of bad-consequence avoidance. These risks are compounded by technology and business shifts that make cybersecurity postures complex, with a backdrop of a cyber skills gap. In addition to securing themselves against these threats, organizations must also document and demonstrate the measures they are taking to meet evolving regulatory and compliance requirements. Because risk is accelerating, governing bodies all over the world are mandating new and increasingly complex risk management processes.

EA: How does the Fortinet solution set address these challenges, in your opinion?
PQ: The Fortinet Security Fabric is an intelligent framework designed around scalable, interconnected security combined with high awareness, actionable threat intelligence, and open API standards for maximum flexibility and integration to protect even the most demanding enterprise environments. Fortinet is the only company with security solutions for network, endpoint, application, data center, cloud, and access, designed to work together as an integrated and collaborative security fabric. This also means we are the only company that can truly provide a powerful, integrated end-to-end security solution across the entire attack surface. Fortinet's security technologies have earned the most independent certifications for security effectiveness and performance in the industry. When woven together, the Fortinet Security Fabric closes gaps left by legacy point products and platforms by providing the broad, powerful, and automated protections today's organizations require across their physical and virtual environments, and from endpoint to the cloud. Today, the Fortinet Security Fabric is a seamless architectural approach to security that is designed to connect security components into a unified, future-proof solution. This vision aligns perfectly to what it takes to deliver automated, intelligent, scalable, and integrated security architecture for today's digital economy.

EA: Fortinet has been developing security tools for many years. What does the company see as the best platform development strategies moving forward?

PQ: The isolated, proprietary security devices most organizations have deployed over the past decade are simply not designed to solve today's cyber security challenges. Data, applications, and transactions traveling between a variety of users and devices often need to span multiple borderless domains and ecosystems. Traditional security solutions, however, tend to operate in isolated security siloes and add complexity to networks striving for simplicity. To support today's dynamic networks, security needs to recognize and understand every device on the network, dynamically segment traffic at the access point based on policy, and monitor and protect data and resources across the entire distributed environment, from IoT, through the network core, and into the cloud. Security is evolving from point solutions to being ubiquitous security everywhere; to having security integrated so that it works as a team; to having that team-oriented security simplified so it doesn't overwhelm the operators; and importantly in the future, to automation, and in fact, so highly automated that it will execute the intent of the operator and security team.

EA: Based on your experience, what do you see as the threat landscape in the coming years?
PQ: The threat landscape is as bad as it's ever been, if not worse. Cybercrime is big business, and is growing at an exponential rate. British insurer Lloyd's of London estimated the cybercrime market at $400 Billion in 2015. Today, just two years later, the World Economic Forum estimates that the total economic cost of cybercrime to currently be $3 trillion. Cybersecurity Ventures is predicting that cybercrime will cost the world in excess of $6 trillion annually by 2021. Cybercriminals capitalize on finding new ways to exploit increasingly complex network environments. To stay ahead of detection technologies, cybercriminals are continually developing new techniques and resources to bypass security and evade detection. Ordinary users and sophisticated businesses alike have inadequate appreciation for what adversaries seek to do to them. Adversaries – particularly nation-state ones – have both the motivation and means to do very harmful things. The way to address those problems of speed and scale, is, in part, by employing solutions like Fortinet's that embrace automation and integration. Those attributes will go a long way toward reducing overall risk.

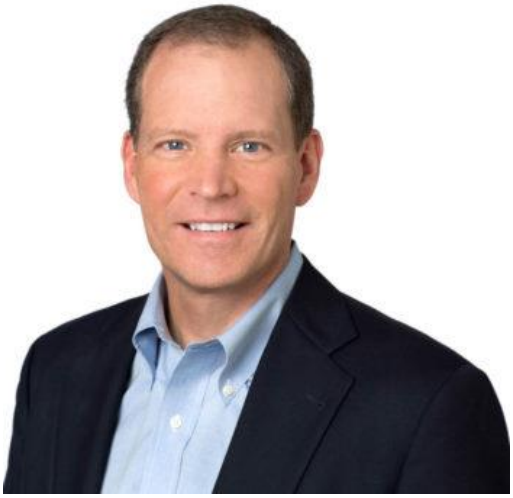

Fusing Data and Cyber Indicators to Support SOC Ops

Detection, hunting, and response benefit from fusion platforms the turn all-source into actionable intelligence

Matt Jones, CEO of E8 Security

Cybersecurity operations have evolved from people leaning forward in their SOC desk and squinting at screen after screen of firewall alarms, to advanced automated platforms processing fused data from multiple sources to generate accurate indicator detection. This evolution goes further: Where incident response was previously based on human-time processes with ad hoc support information, the new approach also makes full use of automation to organize response activities, and help SOC teams understand the process for, and consequences of, cyber response decisions. Matt Jones, CEO of E8 Security, knows quite a bit about building platforms to support SOC operations. He was kind enough to share his insights into how this technology has helped our community, and what we should expect in the future.

EA: What is meant by the term "fusion" in the context of your platform? Is this part of the automation?
MJ: Fusion is more than a name; it is core to our approach to security. Since day one, we have focused on transforming security operations by automating the learning of user and device behaviors to discover malicious activity unknown to security analysts. With this, we aim to make security operations more proactive. The E8 Security Fusion Platform surfaces early warning signs associated with critical cyber threats, such as compromised systems and credentials, privileged access abuse, command and control activity, and lateral movement. The Fusion Platform shows connected behaviors – critical, suspicious and normal – for every user and device. This enables security teams to quickly visualize the relationships between targets and to uncover hidden attack patterns, resulting in accelerated investigations and a more proactive security approach.

EA: Does the SOC threat hunter understand the power that's available in modern platforms to support their mission?

MJ: Absolutely. We at E8 have a simple mantra: "Follow the behavior, find the threat." Behavioral intelligence is a very powerful method of practice and it's one that threat hunters understand. The Fusion Platform starts by fusing all the information about users and their respective devices so that analysts and threat hunters instantly have the context they need to understand what is normal in their environment and investigate suspicious behavior patterns quickly. For security teams to fully understand whether a behavior indicates a threat, they need to see and relate behaviors that are happening on the network, on their endpoints, as well as what their users are doing. The Fusion Platform doesn't just ingest data from network and endpoint sources, it runs analytics on that data, and ties all the behaviors together, providing security analysts and threat hunters alike with holistic insights into their data, users, and devices that they didn't have before.

EA: What sort of features and functions are you seeing being requested by your customers?
MJ: One of the most common challenges for security teams is the complexity of deploying behavioral analytics products. We often encounter customers who are ready to pull out their hair due to the sheer magnitude of time and cost that most behavioral analytics (or UEBA) vendors require. For E8 Security, our approach is to keep security infrastructure simple. Our customers don't need to replace their SIEM, install additional network sensors, or deploy yet another endpoint agent. We make implementing behavioral analytics simple by sitting out-of-band, ingesting all security data from existing sources and technologies already deployed, and easily integrating into the security operations workflow. E8's Fusion Platform is extensible, built for big data, and makes your existing security stack smarter.

EA: Does the transition to hybrid cloud change the nature of the SOC team's mission?
MJ: Part of the Fusion Platform's 'no-fuss deployment' comes from its flexible form factor; it is available as hardware, software, or as a private cloud instance, and fits on top of an organization's existing data lake. None of its threat detection capabilities rely on user-generated or maintained correlation rules or thresholds. This means the platform is not limited to detecting threat activity and indicators that are known and it does not require an army of security professionals to continuously create, review, and rewrite correlation rules as the enterprise and threat landscape changes. With the Fusion Platform, you're identifying behavior in real time and adapting as your business environment evolves.

EA: What are your predictions for the future regarding our industry's ability to detect and stop cyber attacks?
MJ: You must be constantly improving to adapt to our changing world. Operating with this belief, E8 is continuously in the process of building new behavior models that can be applied to our platform, creating new ways to retrieve and analyze supporting data, and addressing compelling problems and use cases. Over the next 12 months, accountability for security risks will be mission critical the for the C-suite. Cyber threats will be scrutinized along with financial and operational risks endangering the company. CIOs and CISOs will need to enhance existing capabilities to detect threats inside their organization. To identify the potential presence of attackers based on activities that are not considered normal and to contain those activities as quickly as possible, security teams need to implement new strategies that incorporate real-time data and machine learning. This operational shift will require self-learning behavioral analytics to detect the early warning signs of today's most

critical cyber threats, such as malicious insiders, external attackers, and targeted malicious software. Our goal is to guide our customers toward a more productive path, bolstering their detection and response. E8 will continue to equip security teams with the visibility and the insight they need to protect the corporate assets from growing threats. At the same time, we're educating the security community on the value of behavioral intelligence, helping security analysts understand how to be more proactive.

Enterprise Solutions for Vulnerability Management

Changing the game for finding and closing cyber security vulnerabilities with accuracy and expediency

Larry Hurtado, President & CEO of Digital Defense, Inc.

For quite some time, vulnerability management meant patching. While effective patch programs are certainly a requirement in any enterprise, this is just a small portion of the types of significant concerns that must be addressed in a modern vulnerability management program. Sadly, however, too many groups are still running insufficient programs in this area, thus leaving their network and applications vulnerable to a range of threats. One might even construct a maturity model for vulnerability management to help build a roadmap for any group still somewhat stuck in patch management. We recently sat down with Larry Hurtado, CEO of Digital Defense to gain insights into platform-based vulnerability management, as well as to discuss the prospects for measuring maturity levels of organizations.

EA: Larry, please start by giving us a good working definition of vulnerability management in the context of enterprise cyber security?
LH: When asked to define vulnerability management, I always like to reference the quote found in NIST SP 800-40 which cites the Benjamin Franklin saying: "An ounce of prevention equals a pound of cure." Patch and vulnerability management is the "ounce of prevention" compared to the "pound of cure" that is incident response. I also like the Wikipedia definition of vulnerability management: "Vulnerability management is the cyclical practice of identifying, classifying, remediating, and mitigating vulnerabilities particularly in software." Why do I like this definition most? Because it notes the cyclical nature of this very important information security program component. With our clients, we always emphasize the importance of managing vulnerabilities on a full lifecycle basis. The management of vulnerabilities is central to modern enterprise cyber security. It involves traditional issues such as tracking and optimizing patch management, but also includes the obligation to maintain an accurately quantified and qualified understanding of

gaps that might exist in security protection and compliance. This is the focus of our team at Digital Defense.

EA: What are the advantages of automating the vulnerability management process?
LH: The advantages include increased processing speed through reduction of manual data processing; improved results accuracy by reducing human error; and enhanced end-to-end program effectiveness resulting in security risk reduction. If we assume that systems used in a vulnerability management program produce accurate results and that the systems are set up and configured properly, then the advantages noted above should be realized. However, there are technical challenges associated with managing vulnerabilities in an IP-based network that have not been addressed by many popular security systems. The challenges are specific to a systems ability to keep track of network devices over time. As the industry continues to head down the path of increased levels of security automation, it is critical organizations understand the importance of selecting vulnerability scanning systems that effectively account for these device moves and changes. Otherwise the organizations' security ecosystems will experience garbage in, garbage out.

EA: If changes are made to infrastructure on a regular basis, does this complicate management of vulnerabilities?
LH: When Digital Defense opened its doors for operation in January 2000, we believed organizations would need to assess network security on a regular basis. Although security assessments were carried out much like financial audits on an annual basis when we first started, we believed the rapid change in the threat landscape would result in the need for more frequent security posture evaluations. Sure enough, before too long, we were asked by our clients to start running assessments more regularly. When this started happening, we noticed something. Network device parameters like IP address, MAC address, and host names were frequently changing on all forms of network devices and not just workstations configured for Dynamic Host Configuration Protocol (DHCP). We also noticed security technologies of all sorts were not effectively accounting for these changes. When one does not properly account for these network device changes, the accuracy of a tool and/or a security ecosystem is impaired. To the unsuspecting user of these flawed tools, the outcome can be a devastating security breach. To those that do understand that these weaknesses exist in security systems and platforms, the outcome is typically a dramatic increase in staffing either to manually track the changes or in the deployment of agent technology. For example, we frequently hear how organizations are employing individuals to export data from certain vulnerability management systems to manipulate the results from these systems in spreadsheets to compensate for this issue to obtain a clear view of what is happening in their security remediation programs. Rather than deploy a vulnerability management platform that requires clients to increase staffing levels, Digital Defense decided to develop technology that accounts for this change automatically. We patented technology enabling us to fingerprint technology more effectively and to keep track of network devices over time, even with significant device parameter changes. This patented technology produces high quality results and eliminates the need for clients to add personnel to manually account for the changes or manage agents.

EA: What's been your team's experience measuring maturity of organizations in their vulnerability management programs?
LH: Our organization produces assessment tools for the "Do-It-Yourselfers" of the security world to use. These tools are geared towards helping organizations perform security assessments and compliance audits. We are also an organization that leverages the tools we produce enabling us to serve in the capacity of a Managed Security Service Provider (MSSP) once again focused on security assessments and compliance audits. In this capacity, we are viewed by our clients as an independent expert, with the primary objective of helping client organizations reduce security risk as rapidly as possible. We perform hundreds of thousands of penetration tests each year and we scan millions of network devices each month. As we perform this testing, we measure the security risk of the organizations we serve using a scoring mechanism called Security Grade Point Average or Security GPA®. We keep track of these scores each year, and then rank the clients we serve based on a composite Security GPA value determined from samples taken each calendar quarter. Each year for the last 9 years, we have recognized the best performing client in one of three categories: Small, medium, and large network. We present trophies to the winning team each year in each of the three categories. This regular measurement of an organization's security risk posture resulted in our wanting to go further in trying to help our clients determine how they can strengthen their vulnerability management programs. As a next step, we've developed a Vulnerability Management Maturity Model (VM3) that includes a questionnaire for self-assessing the degree to which vulnerabilities are properly covered in the enterprise. Organizations with higher VM3 scores tend to be more complete and formal in their VM operations than the ad hoc capabilities, processes, and methods used by teams with lower score. The VM3 tool is available on our website.

EA: Do you ever think we will reach the point where software and systems will be robust and correct enough that we do not have to manage vulnerabilities?
LH: If humans continue to be involved in the development of software and systems, then I would say no. Humans make mistakes as we all know. Whether mistakes are made in the actual writing of code by software developers or afterwards when software applications and systems are deployed into networks by IT personnel and system set-up and configuration mistakes are made, the need for vulnerability management will remain. I must emphasize I am assuming humans remain involved here. Obviously, we are headed down a path where that may not be the case. Artificial Intelligence (AI) may advance to a point where humans will not be required to perform traditional vulnerability management tasks. However, even with the advancements of AI, it seems to me we will need to manage vulnerabilities in the AI systems themselves. Of course, if you agree with Elon Musk and Shane Legg, when we get to the point of needing to determine how best to manage vulnerabilities in AI systems, we will not need to worry about it because humans will no longer exist, right?

Applying Advanced Deep Learning to Cyber Security

Computational advances in neural network processing and algorithms enable new applications to cyber security

Eli David, CTO of Deep Learning

If you had to depict the evolution of malware detection algorithms, you might create an ordered list as follows: Signature-based, behavioral, machine learning, and deep learning. This ordering follows both a chronological sequence, as well as improvements in power and accuracy. What enabled the advance from traditional machine learning to deep learning were community improvements in the underlying neural platform as well as better algorithms. Eli David, CTO of Deep Learning, spent some time with us explaining how this progression to deep learning enables new generation of improved malware detection solutions. This starts with endpoint, but the possibilities for deep learning in cyber security seem endless.

EA: What is the difference between machine learning and deep learning?
ED: Deep learning, also known as deep neural networks, is a sub-field of machine learning, and is inspired by the way our own brains operate. That is, many neurons are connected to each other through many connectors called synapses. These connections learn by being exposed to training data. While traditional machine learning depended on manually extracted features, such as the distance between pupils in facial recognition, deep learning directly operates on raw unprocessed data, without relying on a small list of human-selected features and patterns, and thus, obtains a substantially higher accuracy. In the past few years, deep learning has resulted in the greatest leap in performance in the history of AI and computer science.

EA: Do you need massive computational power to run deep learning algorithms?
ED: We used to require large-scale computational power to train deep learning models. However, during the past few years GPUs (graphical processing units) have been successfully employed for training neural networks by orders of magnitude faster, to the extent that today it is possible to train a deep learning brain on hundreds of millions of samples. This speedup is due to the massive parallelism of GPUs, and the inherent suitability of neural networks for this kind of parallel processing.

EA: How does a deep learning engine learn from live samples? Help us understand how this data is fed to a neural network.
ED: Typically, a large dataset of training data is used. This includes many samples, which, in the case of cyber security, will be computer files. Each sample is labelled, which means that for cyber security, files are marked as being malicious or legitimate. During the training time, these files, which can typically include hundreds of millions of samples, are fed into the deep learning engine, and it gradually optimizes its brain – the synapses – to be able to better separate between the different categories of samples. Obviously, for cyber security, this means differentiating between malicious and legitimate files.

EA: How does your company apply this technology to endpoint solutions?
ED: While other endpoint protection solutions rely heavily on either manually specified rules, or use traditional machine learning based on manually specified features and patterns, we rely on end-to-end deep learning. What this means is that the input into the engine is the raw file, made up of raw byte values, without any preprocessing, and the output produced by it is the classification of whether a file is malicious or legitimate. The training takes place on hundreds of millions of files in our laboratory, running on GPUs. When the model has already learned, we put a copy of this pre-trained brain on each of the endpoint devices (i.e., laptop, desktop, server, mobile) that we protect. During runtime, for any new file on the device, the brain scans it within a few milliseconds, and if it is deemed malicious, it is prevented prior to execution.

EA: What other security applications do you see benefitting from deep learning?
ED: Wherever deep learning has been applied in the past few years, it has yielded huge improvements. We expect deep learning to revolutionize other domains within cybersecurity as well in the upcoming years, including traffic analysis, data leak prevention, phishing detection, etc.

Software Defined Perimeters and Data Center Security

Modern infrastructure security now requires the optimal combination of virtualization and software-defined solutions

Leo Taddeo, CISO of Cyxtera Technologies

The dissolution of the perimeter is the biggest change in enterprise architecture since the invention of firewalls. Sadly, while every enterprise team admits to this clear shift, replacements for the traditional DMZ have been slow to deploy. This stems from a combination of organizational inertia and weak distributed control offerings from most security vendors. Luckily, that is beginning to change, and Leo Taddeo, CISO of Cyxtera agreed to share with us the foundations of Cyxtera, a new cyber security company that offers a range of advanced modern enterprise and infrastructure security solutions, including a powerful approach to software defined perimeters – a clear alternative to the traditional hardware DMZ. Here is what we learned from Leo:

EA: Leo, is a software-defined perimeter easier or harder to implement than the traditional hardware DMZs we are all so familiar with?
LT: On the surface, you would be tempted to say that the distributed nature of a software defined perimeter, combined with the relative unfamiliarity some security experts might have with the approach, would make it considerably harder to implement that traditional hardware DMZ. And furthermore, just leaving your legacy DMZ in place, is obviously less design and implementation effort, although you're probably making up that time with added response costs. But the truth is that an SDP is as straightforward to create as hybrid cloud deployment, which is to say that it is getting easier, cheaper, and more streamlined.

EA: How does remote access work into and out of a software-defined perimeter? Do CISO teams need to procure special hardware to support this?
LT: You're hitting on one of the primary value propositions for Cyxtera, namely the intimate relationship between the SDP and remote access solutions. We've married the advanced software-defined perimeter solution from the Cryptzone team with the virtual segmentation capabilities of Catbird into a new Cyxtera platform that does not require

procurement of new hardware and does not reduce remote access options as an organization gravitates to cloud.

EA: Do you see data center security as being the next Big Thing, with so many companies turning over the reins for their application hosting to the larger cloud service providers with their enormous data center infrastructure?
LT: For companies using virtual data center services in the cloud, the security obligations shift but do not go away. So, for these large cloud infrastructure providers, we do see security as the next Big Thing on their list of obligations to their customers. This is basically true as well for companies virtualizing their own data centers. The good news is that the cost advantages of cloud migration can help offset the costs of additional security.

EA: Cyxtera is an interesting combination of several different component companies. How do you build a common culture from so many underlying creative firms?
LT: The common denominator is excellent in enterprise support using advanced technologies. The synergies are obvious once you begin to integrate the teams and the different capabilities. I've mentioned the synergy between Catbird and Cryptzone already as a prominent example.

EA: Leo, you have a long legacy in cyber security as a law enforcer. Do you see cybercrime speeding up – and can the good guys keep the more serious attacks from occurring?
LT: The question of whether good guys can stop attacks is not the way to look at this. We will never stop attacks. Instead, we must focus on reducing risk, and I do see the good guys closing the gap in that area. Regarding the speed of these attacks, I think you would agree that the time gap between observed attacks in any environment is closing to the point where most groups would claim to be under constant attack. This makes for a challenging risk mitigation environment.

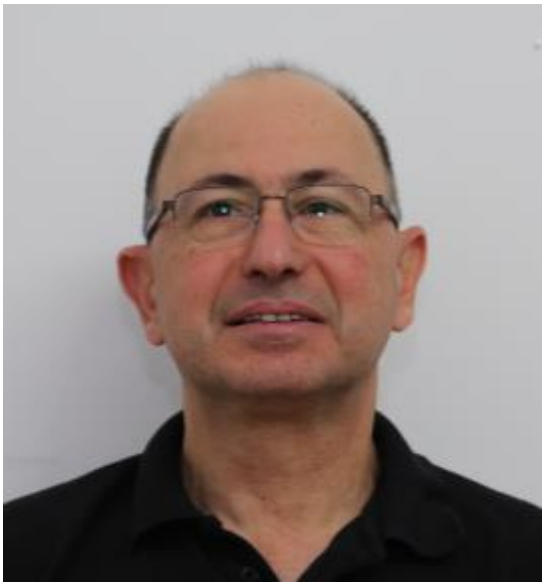

Automated Enterprise Risk Management and Decision Support

Cyber security solutions that enable CISOs to monitor, measure, and manage their organizational risk.

Elon Kaplan, CEO of Cytegic

The essence of effective cyber security management is good decision-making based on sound judgment. Most CISOs recognize, however, that decisions are rarely made with time to reflect, balance options, and carefully consider consequences. Instead, cyber security decisions are selected in the rush of the typical onslaught of incidents, challenges, and risks that characterize the modern enterprise security ecosystem. Elon Kaplan, CEO of Cytegic, sat down with us to explain how his firm develops a platform that assists in this process and brings some calm sanity to the CISO's decision making process.

EA: The Cytegic platform is called Automated Cyber Risk Officer (ACRO), which seems to imply perhaps that it might take the place of the risk official. Is this a correct view?
EK: We like to look at ACRO as a force multiplier for a risk officer faced with a mountain of risk, but only a teaspoon to remove it. ACRO also enables organizations without someone in this role to establish a starting point regarding risk management and mitigation. Essentially, ACRO allows senior decision makers such as the CIO, CFO, business owners, and boards to demystify cyber risk into something that is simple, actionable, quantifiable and translates to dollars. ACRO allows organizations to become proactive and operational regarding the management of cyber risk. Once deployed ACRO can give alerts every time cyber risk exceeds the organizational threshold; it can be used for cyber insurance renewal, premium renegotiation, and coverage; it can be used to establish a dialogue with the CFO to optimize budget against risk appetite; and it can be used at the risk committee level to review and ask tough questions about the preparedness of the organization for the inevitable cyber attacks that will come.

EA: Elon, you have a background in organizational psychology. Did this influence the design of the Cytegic solution?

EK: My unique background has influenced ACRO from the standpoint of communication. I have a deep understanding of the need to translate cyber risk management into terms that can be understood by the organization, including at the executive level. Communication is critical while being different and complex across each silo of an organization. The key is to bridge the gap, and that is what we have done with ACRO. Cyber risk is an executive decision making process, and my expertise allows me to translate the complexity of the cyber world into the perspective of the senior executives, while also leveraging my abilities to the operational level as well. My background in research was the base for some of the very complex statistical algorithms Cytegic employs via ACRO. For example, predicting human behavior which is the most problematic and chaotic to predict was part of the process to translate specific methodologies into our system.

EA: Tell us about how your platform collects information and supports risk-based decision-making.
EK: Using a big-data analytical system, with patented methodology and algorithms, allows ACRO to calculate thousands of correlations between threat vectors and defense vectors at any given time. The outcome of these calculations gives insight into cyber risk per every business asset and business environment across the organization. Moreover, this gives the user a clear action-items list for minimizing risk and focusing funds only where it matters. By using ACRO, users save time and money, but more importantly they utilize their current resources to lower risk and financial impact to their most valuable assets. On top of this, the solution allows users to perform almost daily compliance as compared to a periodic one. The threat landscape is correlated utilizing open sources and any third-party threat intelligence vendor.

EA: How does your platform integrate with familiar risk management processes in place in most companies?
EK: ACRO can collaborate with any standard, framework, security, or risk-related technology. ACRO can integrate data and processes to cater to any level of customization that each unique organization may have. ACRO collects and synthesizing the control maturity indicators, creating with a click of a button, dedicated dashboards and reports per each standard, such as ISO 27001, NIST, and HIPAA. Internally, ACRO automates the collection of maturity indicators to minimize subjective human error and takes a data-driven approach to control maturity assessment – a process that up until now was done by hand. The quick and simple, yet comprehensive, simulation capability allows users to create what-if scenarios to plan and assess the impact of future changes to control deployment or threat landscape before taking a step. This allows for smart budget planning and road-mapping.

EA: Share with us the process of how analytics are used to synthesize inputs to produce useful recommendations for CISOs.
EK: ACRO uses a unique patented set of algorithms that correlate between internal security control maturity score and external threat landscape trends and patterns. The system works seamlessly to present the user with the calculated risk scores, alongside financial impact analysis and "what-if" scenario analysis. Analytics is part of every step of the process – from collections and processing of raw material, through big-data trend analysis

and risk management. The main added value from the system is an importance score for each security control that is basically a call for action based on the impact each control has for minimizing the risk. The system is basically telling the user – after calculating thousands of correlations, that these are the most important places to allocate resources to mitigate the most threatening risks. By doing all the above, ACRO acts as the CISO's best support platform, taking on the heavy duty load, and allowing for data-driven decision making.

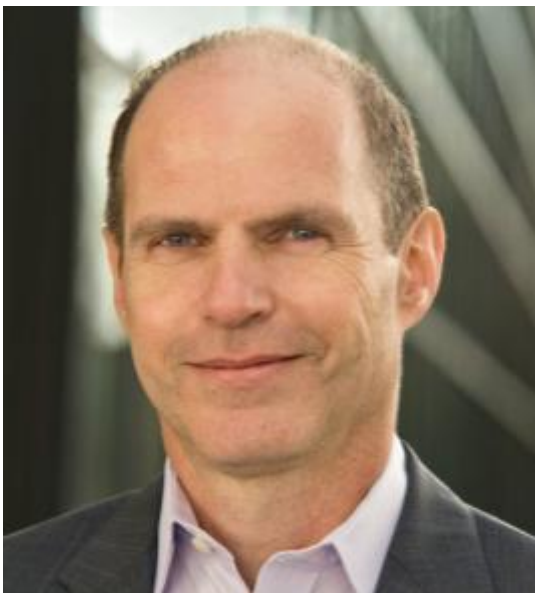

Eliminating Cyber Threats Using Insider Privileges

Most advanced offensive actors will readily admit that access to privileged accounts is a critically important attack vector

Udi Mokady, Founder, Chairman and CEO of CyberArk

Enterprise security has traditionally focused on investing in perimeter-based security solutions to protect networks. While such work remains vital to the secure operation of any business, experience suggests that motivated external attackers will always find a way in. Privileged accounts and credentials offer a lopsided advantage to offensive actors trying to exploit vulnerabilities in an enterprise. Once attackers break through the perimeter and compromise privileged credentials, they become an insider with the ability to move throughout the network, virtually undetected. Udi Mokady, Founder, Chairman and CEO of CyberArk, understands this risk well, and he was kind enough to share his expertise with us regarding the best way for protecting privileged accounts.

EA: What is the primary difference between privileged and non-privileged accounts?
UM: Privileged accounts are everywhere. They are in every networked device, database, application, server, and social media account; they are on-premises, in the cloud, and in ICS systems. This explains why privileged accounts are often referred to as the keys to the IT kingdom. They provide administrative access to business-critical applications, systems, and networks in an organization. They trace their lineage from the early root accesses made available to system administrators of operating systems and are foundational to administering IT and running the business on-premise and in the cloud. The community has come to recognize that privileged accounts are the preferred means by which insiders and external attackers gain power, and are able to assert control over a network. Regardless of where the attackers start, they need privileged credentials to move throughout the network. Both internal and external attackers look the same once they have compromised privileged credentials. While managing all levels of IT access is important, the consequences of privileged account exploitation can be severe, which is why protecting them must be a priority.

EA: Is it easy for a typical business to take inventory of their privileged accounts?
UM: The first step for an effective risk management program is to quickly identify privileged accounts wherever they may exist across the enterprise. This can be a challenge for some organizations because of the sheer volume of privileged accounts that exist across the enterprise including user accounts, SSH keys, service accounts, devices, and applications. The CyberArk Discovery and Audit (DNA) tool is one way organizations can easily identity these accounts and quantify security risk within enterprise networks. By better understanding the size and magnitude of their privileged account security risk, organizations can more effectively build a business case for a privileged account security program.

EA: What techniques do you use at CyberArk to protect these highly-privileged accounts?
UM: Most organizations also don't fully understand that privileged accounts are used in virtually every cyber attack, so deploying privileged account security needs to be one of the very first steps an organization takes to secure its systems. Securing privileged accounts is also the first action organizations take following a breach. We provide organizations with an easy-to-use methodology, which we refer to as the "30 Day Sprint," to prioritize the implementation of controls for protecting privileged credentials. Once organizations have identified where privileged accounts exist in their enterprise, they must prioritize and give precedence to the riskiest accounts. This means implementing controls on the most powerful accounts first, such as domain administrator accounts and administrator accounts with access to large numbers of machines, as well as application accounts that use domain administrator privileges. We advise customers to be realistic about addressing the volume of accounts, they don't have to boil the ocean to achieve quick wins and demonstrate tangible results. Organizations should work quickly to get initial controls in place and make improvements over time. For example, accounts for workstation users should not have administrative privileges, but breach survivors say this is one of the more difficult practices to implement and maintain due to the sheer volume of workstations.

EA: How does a digital vault work? Does it create a single point of attack for the bad actors?
UM: CyberArk was founded to help organizations build a security strategy from the inside, focusing on locking down the keys to the IT kingdom. This is how the concept of digital vaults and privileged account security was created. At the core of the CyberArk Privileged Account Security Solution is the CyberArk Digital Vault, which contains a highly secure repository, behind multiple layers of security, which stores privileged account credentials, access control policies, credential management policies and audit information. CyberArk is first and foremost a security company, and we design our products with a "security first" mindset. The Digital Vault software is intentionally designed to minimize the attack surface and maximize the security of privilege account information. In addition to internal vetting and testing, CyberArk also submits its products to external organizations for independent testing and security validation. Through this process, the CyberArk Privileged Account Security Solution has achieved ISO 9001, Common Criteria and United States Department of Defense UC APL certifications.

EA: Do you see privileged account security protections becoming more uniformly applied across different systems and applications?
UM: Today, many organizations still underestimate the scope of the attack surface that privileged accounts create. It's not unusual for larger organizations to have hundreds of thousands of privileged accounts. That attack surface is expanding exponentially as organizations migrate to the cloud and invest in new DevOps and endpoint technologies. While our business was initially driven by organizations in highly regulated industries, with greater recognition of the risks posed by privileged accounts, privileged account security has evolved from an audit and compliance solution, to become a critical layer of IT security, and essential to every organization's risk management strategies. We view the privileged account security market as a green field opportunity. This is because virtually every organization runs on technology, which requires protecting the privileged accounts that control that technology. If they lose control of their technology, they effectively lose control of their business.

Integrating Cyber Attack Detection and Mobile Security

Eliminating cyber attacks throughout the enterprise including those targeting mobility and cloud assets

Kirsten Bay, CEO of Cyber adAPT

Traditional attack detection capability focused on conventional cyber threats that originated on some malicious actor's PC, traversed a progression of local and wide area IP networks, and then targeted a valued asset on some server. While this use-case remains valid, the need has clearly arisen to take mobility and cloud into account in any attack detection platform. This implies understanding how to detect indicators in mobility management, cloud applications, and other modern enterprise infrastructure. We had the opportunity to connect recently with Kirsten Bay, CEO of Cyber adAPT, to learn more about how attack detection and mobility can come together in an enterprise security solution.

EA: Kirsten, what are the essential elements of a successful, modern attack detection platform?
KB: In a simple sentence, the ability to deliver relevant context to users such that they can rapidly address events that they understand have potential impacts on their environment. This means that the platform must be able to take in a variety of data inputs such as host and device type, threat intelligence, user behavior, analytics and classification of data form and origin, and packet analysis, just to name a few. While we in the security industry have made great strides in developing ingest engines and platforms that utilize advanced analytics to deliver this type of platform, we still have work to do in helping incident responders understand the relevancy of these attacks. In other words, providing the "what does it mean to me" element is the truest form of delivering a successful attack detection platform since not all attacks are created equally.

EA: How important is it for platforms to utilize advanced analytics to reduce cyber risk?
KB: Advanced analytics has become an important element for dealing with the large data sets that are derived from a platform's ability to capture a granular level of data. That said, one of the challenges is that this level of data, even with advanced analytics, can result in too many false positives. This explains why we believe that the integration of deep

intelligence data combined with advanced analytics creates a higher level of fidelity and a lower rate of false positives.

EA: What's been your team's experience integrating mobility into your attack detection platform?
KB: We have found that it provides a significantly improved level of protection for our customers, as well as detection of mobility-related events related to individuals as well as the mobile device. The protection feature keeps users and their data from being compromised by man-in-the-middle attacks while securing data in motion, and our attack detection solution, including for mobile, inspects traffic for attacks before it passes through the firewall to ensure we are catching bad events on the edge.

EA: Do you see the biggest risks emerging from mobility in the coming years?
KB: Absolutely. In basic terms, users are defaulting to their mobile devices more and more to complete work activities, and that will ensure that attacks will continue to be sourced and targeted at the device level. In broader terms, we are already seeing this trend rapidly grow with the explosion of IoT devices that are the genesis of many types of attacks. From our standpoint, a mobile device is not just a phone. It is any device that operates off-premise that connects to a network. I often hear people say that they haven't seen many attacks resulting from a mobile device, and that therefore, the risk is insignificant. I challenge that point of view when anyone can click on a phishing email on a device that promotes lateral movement by the adversary into a network, because these devices allow a multitude of vectors that didn't exist before.

EA: Kirsten, you are an amazing role model for youngsters interested in cyber security, especially young women. What advice would you have for them regarding our industry and successful careers in cyber?
KB: I would advise to be as curious as possible. Think about how economies and societies are impacted by both technology and the risks posed by the speed at which we want to do everything. Be curious about the connection and intersection points beyond how technology functions. Cyber security has a large growth curve, but we do need to ensure a multidisciplinary approach to how we solve for this growing threat. I am also very hopeful that the generation entering the workforce will work to be more inclusive and supportive of each other. I am dismayed by much of the news in the technology world at large that makes us seem unaccepting and closed. The only way for us to continue to excel as leaders and innovators is to continually seek new ideas and perspectives. Those of us leading now must work to change that perception, and the younger generation must demand it.

Advanced Endpoint Cyber Solutions for Stopping Breaches

Using advanced algorithms combined with cloud-based threat intelligence provides best-in-class protection

Dmitri Alperovitch, CTO of CrowdStrike

Endpoint protection has always been a primary concern for enterprise security teams, given their prime targeting by malicious actors. Despite this emphasis and the clear availability of many different endpoint security options, the risk associated with enterprise PCs, and now mobiles, tablets, and IoT devices, continues to rise. One problem is the over-reliance of traditional endpoint protections on signature-based antivirus software, which has been discredited as ineffective against variants. The good news, however, is that modern vendors have improved their technical approaches considerably, using the best available algorithms and advanced methods to combat endpoint risk. Dmitri Alperovitch, co-founder and CTO of CrowdStrike, knows a thing or two about this problem, and sat down with us to share his unique perspective on cyber risk management for endpoints.

EA: Your firm claims to center on the theme of 'stopping breaches.' Is this really a tractable goal?
DA: At CrowdStrike, we like to say that we don't *have* a mission, but that we are *on* a mission to stop breaches and adversaries. While most security companies have built solutions to protect against malware, exploits, malicious websites, and unpatched vulnerabilities, there is a fundamental flaw with this approach: Any malware-centric defense leaves organizations vulnerable to attacks that don't leverage malware. In reality, malware is responsible for only 40 percent of breaches and advanced attackers are increasingly leveraging malware-free intrusion approaches to blend in and fly under the radar. The CrowdStrike platform provides a solution that protects companies against malware and non-malware based attacks, effectively stopping the breach. The Falcon platform also has one the biggest threat telemetry footprints in the industry, ingesting over 55 billion events per day from millions of endpoint agents deployed in 176 countries.

EA: What are the best available methods for reducing risk on the endpoint?
DA: Today's malware and malware-free intrusions require a comprehensive approach for detection and prevention. The CrowdStrike Falcon platform extensively leverages machine

learning for identification of both known and previously unknown malware files, as well as malicious behaviors. In addition, our IOA (Indicator of Attack) methodology is vital to identifying attacks across the kill chain without having to rely on signatures and IOCs (Indicators of Compromise). This approach applies link analysis and graph traversal technology to determine the intent of an execution action and stop it before the harm is done. You know that someone, even if they are previously unknown to you, is robbing a bank if you see them walking into the bank branch, getting into the vault, and walking out with the money. Similarly, IOAs can identify exploitation activity, privilege escalation, credential theft, lateral movement, and actions on objectives stages of the kill chain, purely by analyzing the intent of the execution activity being observed, without having to know in advance the exact technique or code that is in use. Layering the IOA and machine learning approaches together and combining them with one of the most comprehensive threat intelligence repositories in the industry has been a winning combination for CrowdStrike in stopping over 300 breaches a week across our global customer base. On top of that, we have OverWatch, the best hunting team in the industry, providing a human backup capability on top of the advanced analytics performed by the platform. This team is hunting for adversaries, investigating intrusions, and helping customers contain incidents on a 24/7 basis across millions of machines being protected by CrowdStrike Falcon.

EA: Is cloud-based threat intelligence a vital component of any modern cyber security solution?
DA: I believe it is. Harnessing the power of the cloud is at the heart of CrowdStrike's Falcon platform, and it's demonstrated that the model is about more than just cost, scale, or time savings. The cloud improves Falcon's capability and effectiveness in threat protection. As Falcon ingests data from customers' endpoints, the cloud enables the company to crowdsource protection across its entire customer community within seconds. CrowdStrike's cloud-based architecture offers a level of scalability and a speed of response that is truly vital to modernized security. When the company was founded in 2011, cloud-based endpoint security basically did not exist, as the cloud was regarded too risky. We set on a journey to bring disruption to the industry and enable customers to have better security posture through the agility and crowdsourcing benefits of the cloud.

EA: Tell us about current trends in active response to serious threats. Do most organizations have improved processes for dealing with serious incidents?
DA: It's important for organizations to be proactive and implement a comprehensive security strategy before they get breached. It's never ideal to look back and think "what if?" All organizations should do assessments to understand how prepared they are to respond to an incident. CrowdStrike Services team help businesses answer the question: "Is my security and incident response plan mature enough for the threat environment I face?" Tabletop exercises are one way to simulate an attack to give key stakeholders with the organization exposure to what a real incident may look like and help prepare everyone for the experience.

EA: Given your experience and expertise in cyber, can you share your views and predictions on cyber risk trends?

DA: The biggest risk I worry about today is the escalation in destructive attacks that companies face from both enterprise ransomware and wiper malware. While the incidents of companies going dark due to a cyber incident used to be extremely rare, they are starting to occur with an alarming frequency and with impact measured in hundreds of millions of dollars. Given how unprepared most organizations are at preventing this activity, we are almost certain to see this trend escalate exponentially in the coming months.

Automated Cyber Protections for DevOps and Cloud

The synergy between DevOps and cloud allows for creative cyber security solutions that operate at Agile speed

Carson Sweet, Chief Strategy Officer of Cloud Passage

DevOps is all about reducing the time between customers requesting a capability and developers delivering an implementation. Traditional software engineers might have seen months or even years pass between these two interactions, but modern technology users cannot accept such delays. The result has been a new methodology focused on daily and weekly projects that are Agile, flexible, and designed to optimize communications between users and developers. The security issues that emerge in such environments are intense, and the ubiquity and convenience of cloud infrastructure play a significant role in this new development process. Carson Sweet, Chief Strategy Officer of CloudPassage, is one of the world's leading experts in cloud security and how such protection can optimize DevOps. We asked him recently to share some thoughts on this area.

EA: Are all development teams following some sort of DevOps process today?
CS: I think that is a safe assumption. The days when a customer would be willing to wait months or even years for a system to be delivered are long gone, and the DevOps process fills this gap. I think this reduction in time between customers expressing their need and developers providing functionality is the greatest advantage of the new paradigm.

EA: How does cloud play into DevOps and what are some of the security issues that arise?
CS: Cloud services and infrastructure are perfectly suited to the needs of the modern DevOps team. Just consider the challenge of provisioning servers; in the older method, new functions requiring underlying server support would have to submit to a hardware procurement and provisioning process that could take days or even weeks. With cloud, this is an on-demand, point-and-click operation. The primary security issues, as we've learned at CloudPassage through years of experience, come in several flavors. First, CISO teams need to understand the compliance controls embedded in any cloud environment. Second, they must introduce functionality to protect workloads from attacks. And third, they need to use hybrid cloud transition as a means for reducing dependence on a perimeter.

EA: Do you see virtualized separation such as micro-segmentation as an important mechanism in securing modern software processes?
CS: The closer you can bring security controls to the assets being protected, the better. Micro-segmentation allows system designers to connect a customized protection suite to virtual machines, containers, and workloads in cloud. The virtual separation that results allows for shared secure use of cloud infrastructure and is also an important component of compliance assurance.

EA: Have compliance managers and regulators figured out DevOps and cloud security yet?
CS: They are starting to realize that cloud is more part of the solution than the problem. In the early days of cloud, compliance managers and regulators were correct to be concerned that improper data operations might reside behind a cloud service. Today, however, CISO teams demand evidence and the cloud providers have followed suit. AWS, for example, includes mandatory controls around functions like logging that are going to be improvements on most existing approaches.

EA: What do you see as the major trends in this area in the coming years?
CS: That's an easy question to answer, because the transition to distributed, virtualized cloud infrastructure is already well-underway. We believe this will continue, and that adjacent industries such as industrial control and IoT will follow suit. Mobility will play an important role in this march through hybrid to full cloud adoption.

Using Telemetry to Secure Software Applications

Collecting real-time data from software application run-time environments supports advanced cyber protection

Sameer Malhotra, CEO of CIX Software

Most modern industrial control devices are set up to provide telemetry to management centers for monitoring, calibration, and tasking. It should thus not come as a great surprise that modern software applications can benefit from the same treatment. Specifically, the run-time environment – whether on-premise in the data center or cloud operating system-resident – can offer sufficiently detailed telemetry to allow for high assurance security support for running applications. This is a great advance from the early focus on application scanning and code review that would miss any dynamic threat that might arise during execution. Sameer Malhotra, CEO of CIX Software spent time with us recently to explain how this approach to application security can work.

EA: Sameer, what are the more traditional approaches to application security and what have been the challenges?
SM: It's been our experience that protecting software applications properly for both security and compliance has been a nagging problem in both business and government. Traditional static analysis and application scans, for example, offer some useful assistance, but they are time-consuming and completely miss all run-time threats. Our BUSHIDO product fills this gap by offering comprehensive visibility into applications based on over 115 different parameters. This visibility is essential to preventing advanced threats to applications.

EA: How does a focus on run-time execution and telemetry work for application security?
SM: The optimal protection approach integrates run-time visibility, machine-learned profiles, behavioral analytics, and workflow-driven response. All these elements are fed by accurate application telemetry, which then enables the types of security prevention, mitigation, and response required by CISO teams. This includes visualization, forensic history analysis, DevOps support, real-time behavioral analysis, network lockdown, and so

on. These are powerful application support protections and they are enabled by the embedded protections offered by our BUSHIDO solution.

EA: What algorithms work best to identify anomalies in application behavior? Are they signature-based?
SM: Some signature-based processing continues to be helpful for legacy type threats, but most detection now relies on more advanced machine learning based on profiles. Developing these accurate profiles relies on understanding application and role context, interface requirements, and run-time environment characteristics. These elements leverage behavior analytic methods that process complex events in the context of local security policy. These algorithmic concepts are at the root of the BUSHIDO Product.

EA: Many observers in our industry like to point out that application security is the toughest aspect of cyber security. Do you agree with this claim?
SM: I do agree that application security is the toughest aspect of cyber, and part of the problem is organizational. Both IT and security teams tend to have responsibility for applications, so it's not always clear who should take the lead on protection. At the technical level, applications have such varying functionality that uniform solutions are hard to come by. That's why with BUSHIDO we focus on run-time visibility based on automated platform support.

EA: Can simple network segmentation solve application security?
SM: Simple network segmentation can help establish application boundaries, but it is important to understand normal business traffic patterns with respect to timing and volume to better establish the right parameters for network segmentation.

EA: What other factors do you think drive application security?
SM: You must understand the basics from the hardening status of the OS on the workloads to establishing a good known state of runtime processes with identity in context. Application security means establishing controls across multiple factors such as network connectivity, software, identity behavior, software and OS vulnerability and security policy compliance.

EA: Do you see compliance auditors and regulatory officials driving the benefits application security?
SM: Our early customers are large financial and payment processing organizations, who are consequently subject to compliance pressure. Financial regulators and Industry consortiums such as SWIFT are now driving the focus on protecting the vital crown jewel applications in these organizations. We are now seeing an uptick on CISOs being included into attestations around securing the application environments.

EA: How hard is it for enterprise CISO teams to deploy application behavior analytics and build up to the automated response?
SM: Deployment of BUSHIDO can be completed in hours. We provide a model for both agent and agentless data ingestion. This first step enables the enterprise to get real-time

visibility into the application environment and provides detailed understanding of application dependencies. The next step is to create the application profiles with our simple, automated process which uses machine learning to develop application context white lists for network, process and identity. From there CISO and application teams start to get behavior analytics which they can use to instantiate the workflow driven automated response capabilities on an application by application basis. We provide 50+ rules out of the box.

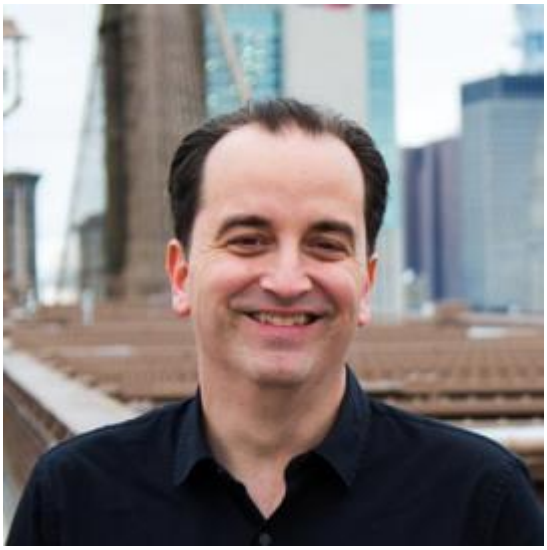

Container-Aware, Real-Time Security for Linux

Advanced threat protection for Linux has been under-served, which is an issue given its pervasive deployment

John Viega, CEO of Capsule8

Cyber security companies emerge from stealth every week, but few were more welcome to see than the recent unveiling of Capsule8, a Brooklyn-based cyber tech firm working hard to help us protect our Linux deployments. It is a well-kept secret that most server infrastructure, including in public clouds, depend on Linux for computing support. While this is good news for expert administrators with strong Linux backgrounds, it has the odd and unexpected implication that many of the commercially available tools to protect servers are not applicable. John Viega, CEO of Capsule8, helped us understand this situation, and explained the technical underpinnings of his team's container-aware security solution for Linux.

EA: John, what are the statistics around Linux use in the data center and cloud?
JV: According to The Cloud Market, more than 92% of Amazon EC2 instances run Linux. About 18 months ago, Microsoft announced that about 1 in 3 Azure instances runs on Linux, and we've heard people claim that this number is now close to 1 in 2. With such widespread adoption of Linux in production environments, it's surprising that the best practice for attack protection for enterprise Linux is stuck in the early 2000's. Our team at Capsule8 is focused on bridging that gap.

EA: What problems do most people face with Linux security?
JV: People tell us that it's difficult to collect and analyze the right data efficiently and easily, without risking bottlenecks or reliability. This is even more true when deployments leverage micro-services. Most people find out about breaches hours or days later, if at all. The world we're enabling allows detecting attacks in progress, and automatically shutting them down as they're happening, without negatively impacting production systems.

EA: Does the approach work differently for legacy Linux deployments as ones newly deployed using your technology?

JV: Our data collection doesn't much care if you're running in a cloud-native environment or a legacy environment. The analytics and automatic attack response can be done a bit better in a more modern environment. For instance, in a modern environment, some pieces of an environment are typically "stateless," meaning they can go up and down without impacting the application. We can leverage that knowledge to provide both more accurate protection, because we have more information on the types of things that *shouldn't* happen. But we can also be a lot more liberal about automated response in a stateless environment (if you're worried a container might be compromised, then often you can just kill it and spin up another one). In a stateful workload, you must be more careful, but can still do things like kill risky connections and alert an investigator. The key here is acting in real time, before any damage is done.

EA: Give us a summary of how security solutions such as yours might prevent unknown threats from occurring.
JV: While there are shockingly many software vulnerabilities, there are far fewer exploitation techniques. We focus on detecting attempts at exploitation, and then, as a fallback, evidence of compromise. For instance, memory-based exploits often involve making non-executable parts of application memory executable. Detecting such things can be highly effective. But if an attack is truly pioneering, we may still notice, for instance, an interactive shell being spawned by a web server, which is a sure sign of exploitation in most places.

EA: You've been at this security game for some time. What are some offensive and defensive trends you are seeing?
JV: One of the most important trends is common to both sides: Automation. Attackers seem to be better at the automation, but the security industry is getting the message. The current number of experienced cybersecurity professionals is not nearly enough to satisfy our cybersecurity needs in the public or private sectors. In fact, ISACA predicts that there will be a global shortage of two million cyber security professionals by 2019. So, the industry can't rely on human expertise alone for protection; we need automation. We are starting to see some leading-edge organizations automate at least part of their response process wherever possible. Capsule8 is designed to not only automate attack response, but to integrate via API to any incident response automation companies have deployed (e.g., Demisto or Phantom). Automation is a huge focus for us because without it, we can't possibly keep up with the growing pressure from our adversaries.

Protecting Endpoints with Micro-Virtualization

Advanced tools are available for containing and separating malicious execution from local assets on an endpoint

Simon Crosby, CTO of Bromium

Micro-virtualization on the endpoint is easy to conceptualize: Containers are used to ensure that if potentially malicious activity occurs, it will hit the walls of a micro-virtualized container, thus protecting local assets on the PC. This concept led to several early product attempts that made security teams happy, but that also displayed the inevitable growing pains of any powerful new technology for enterprise use. For example, keeping track of when-or-how, this-or-that software package was installed before-or-after, this-or-that container was a typical sort of early permuted concern. But the great news is that in the past few years, many of these administrative issues have been dramatically improved. Micro-virtualization is now becoming a conventional control that is indispensable for many CISO teams. Simon Crosby, CTO of Bromium, is one of the pioneers in this area, and he answered some of our questions about how this control can be best applied in a modern enterprise.

EA: Simon, give us a brief explanation of how micro-virtualization works.
SC: Our micro-virtualization technology is based on the Bromium Microvisor, which is a security-oriented hypervisor based on Xen that integrates with the underlying chip hardware to support strong separation of tasks. The resulting hardware-isolated virtual machines support the types of isolation required to create high levels of security in the presence of endpoint malware.

EA: Has it become easier to maintain such capability across many enterprise endpoints?
SC: The extensive deployments we've done in recent years has allowed us to come down the experience curve significantly for endpoint provisioning. Our customers experience much shorter time-to-value and hence see lower protection costs for endpoints. The live intelligence and remediation from Bromium also make it easier to maintain and use the technology.

EA: Does your own solution include any advanced technologies to more accurately detect the presence of malware?
SC: Yes, the Bromium sensor network for detection and response on endpoints combines with our isolation technology to accurately detect malware. Because the isolation offers clear visibility, we sometimes refer to the detection of malware on endpoints as occurring in high fidelity.

EA: What is the role of cloud in micro-virtualized security?
SC: We like to think of our protection solution as enabling endpoints for safe and secure access to cloud, so we certainly play a role in that sense. But our live threat intelligence originates in the cloud and offers enterprise teams the best available information on malware behaviors and patterns in real-time.

EA: Do you think this technology can make a real difference in the prevention of advanced attacks such as APT?
SC: Obviously, by separating potential malware execution from the endpoint resource, the first step in the attack chain for most APTs is addressed. Advanced actors have many offensive methods at their disposal, so no single solution will stop a determined nation-state, but we are confident that our technology plays a vital role in the reduction of risk for APTs and other enterprise cyber threats.

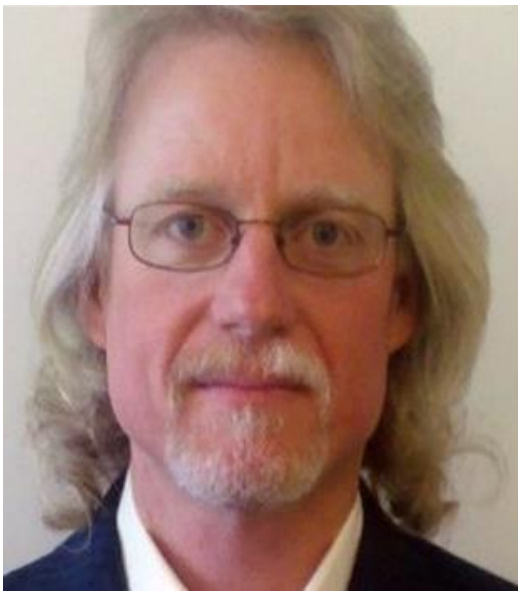

Micro-Segmenting Data Centers and Networks Using Strong Separation and Abstraction

Using embedded trusted identity tokens to enhance enterprise security and policy enforcement at the network protocol layer.

John Hayes, Founder and CTO of BlackRidge Technology

Security policies are typically designed with the intent to check credentials associated with access requests before entry is permitted. The problem is that the network protocols such as TCP are bidirectional and must allow multi-step back-and-forth handshakes between clients and servers to establish identity credentials at the application layer. This violation of most enterprise policies can be solved through advanced separation methods that employ embedded credentials into the network protocol. John Hayes, CTO of BlackRidge, caught up with us recently and helped us understand how BlackRidge supports such separation to truly enforce enterprise policies.

EA: John, explain in a nutshell, how your technology works?
JH: Our technology inserts a cryptographically-secured identity token into the first packet of every TCP session. Across the network, this identity token is recognized and access to network resources is allowed or denied. This allows network and cloud resources to have the identity of the user or device connecting to them before establishing the TCP session. It's really a secure version of Caller-Id for the Internet.

EA: Do clients or servers have to modify their TCP stacks to use your solution?
JH: No, we do not modify the TCP stack. Our endpoint software operates as a shim below the TCP stack and above the device driver. The native TCP stack is unaware of our presence and we operate transparently to the stack and the applications. In this way, we can add identity and authentication to legacy networks and applications without a forklift upgrade of the infrastructure.

EA: Where does a security team position the BlackRidge gateway?
JH: There are several common use-cases and deployment models that position BlackRidge gateways, our identity recognition and policy enforcement points, in different locations. These are positioned at the perimeter of an enterprise to identify and authenticate all external traffic; within the enterprise to provide micro-segmentation; in front of cloud resources to protect those cloud resources from discovery and access, even when using public cloud infrastructure; and in front of the management plane to separate and isolate critical management infrastructure and authenticate access.

EA: Do you support virtualization and cloud environments?
JH: Yes, we support the leading virtual and cloud environments. Public cloud infrastructure does not provide the same discovery protection that traditional physical infrastructure provides, and segmentation within and across heterogeneous environments is difficult to achieve and prove. Cloud resources protected and segmented by BlackRidge do not respond to network scans and network reconnaissance, restoring the discovery privacy and compliance controls previously enjoyed only by physical data centers.

EA: What are the risks of not deploying strong separation solutions such as yours in the data center or enterprise?
JH: Traditional network management relies on using addresses and topology. Addresses can be spoofed and being topologically dependent requires constant synchronization with how the network is currently connected. This creates quite a few hassles today, trying to manage firewall rules and router ACLs. By introducing strong separation via cryptographically-secured identity to the network, BlackRidge provides both authenticated access control on a per TCP session basis and provides attribution information gleaned from our identity tokens to SIEM and analytics systems. This authenticated attribution information is unavailable in traditional data centers and enterprise deployments otherwise.

EA: What special advantages does BlackRidge have for micro-segmentation over other vendor's approaches?
JH: BlackRidge performs First Packet Authentication. This is the ability to determine the identity of the originator of a TCP session on the very first packet of the TCP session, before any response is made to the requestor. This blocks port scanning with no packet leakage problems common to other firewall and application firewall security solutions. By operating at the TCP layer on the first packet, BlackRidge enforces policy at the earliest possible time to provide strong separation with attribution, supports multi-vendor and heterogeneous data center and cloud environments, and provides automation and abstraction from the network.

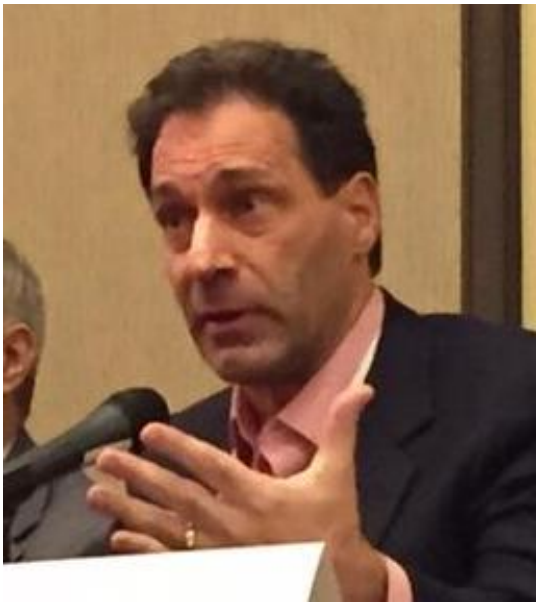

Securing the Industrial Internet of Things

Solutions for protecting ICS service in OT/IT ecosystems requires domain expertise as well as cyber capability

Francis Cianfrocca, Chairman of Bayshore Networks

Securing operational technology (OT) is one of the more challenging aspects of the cyber industry. The industrial control systems (ICS) so prevalent in OT environments incorporate a wide range of technologies from traditional computing to electromechanical and analog systems. Over the past decade, society has become increasingly dependent on the digitization of OT infrastructure. The integration of OT and IT, and the connection OT environments to the Internet offers significant advantages. Data analytics and the resulting Industrial Internet promise to increase productivity, reduce emissions, improve safety, and more. But this interconnection comes at a cost: Increased cyber risk. Hackers now target critical infrastructure such as power plants, public utilities, and factories. Francis Cianfrocca, Chairman of Bayshore Networks, has been working in this area for many years, and is a recognized industry expert. We asked him to share his thoughts on where ICS/OT security is likely to go in the coming years.

EA: Francis, which do you think is harder to protect: OT or IT?
FC: I think you would agree that they are both hard to protect! What makes the protection of OT, ICS, and IoT so challenging is that these infrastructures combine traditional IT systems with the variety of operational, mechanical, analog, and electronic systems that are now connected to the industrial Internet and across IoT networks.

EA: Many pundits such as Ted Koppel have written some scary things about the possibility of attacks on OT infrastructure. Do you agree with such dire predictions?
FC: There are obviously some tough scenarios that we would all like to avoid, and I give Ted Koppel credit for helping to bring more attention to this area. But there are quite a few people, processes, and infrastructure controls in place to mediate cascading OT attacks. The solutions we offer at Bayshore Networks are designed to enforce policy and can reduce risk in a meaningful way.

EA: Is domain expertise required to develop the types of solutions you build at Bayshore Networks?
FC: It is important for our team, but our customers are already experts in their respective domains. The challenge, as I see it, involves combining an IT/OT security team that includes balanced backgrounds in conventional IT, domain-specific OT, and cyber security. Most universities are not teaching these disciplines as a combined entity, and we'd like to see that change.

EA: How big is the problem of non-standard, legacy infrastructure in the protection of ICS/OT?
FC: It's a challenge, and the more legacy some equipment or system, the more difficult it will be to deploy cyber security solutions. In the best case, the legacy systems can be replaced, but in most cases, they cannot. Our approach at Bayshore Networks is to focus on both situations, knowing full well that our customers in oil and gas, energy, government, and other OT sectors need a platform that can provide domain-specific visibility, real-time detection and elimination of threats, and policy-controlled external interconnections regardless of the underlying infrastructure components.

EA: Do you see more industrial engineers developing cyber security skills in the future? (And vice versa?)
FC: As I mentioned above, we'd like to see an integrated cyber curriculum that aligns with the reality of physical and logical security. More cyber security engineers will need OT expertise in the most important domain areas. As malicious attacks move toward that critical OT/IT interface, it will soon not be enough to just understand the IT side of the equation. As a result, industrial engineers will need to learn more about cyber, and security engineers will need to learn more about industrial systems.

Advanced Cyber Risk Reduction via Dynamic Deception

Tools that deceive, detect, and defend systems and networks from malicious attacks are now becoming essential.

Tushar Kothari, CEO of Attivo Networks

Deception has been underutilized in the cyber security industry. This is surprising given the success military organizations have had using stealth, deceptive methods to create uncertainty and confusion in an adversary. Catching an unsuspecting malicious entity is also a great possibility when the use of deception is properly deployed. Tushar Kothari, CEO of Attivo Network was kind enough to share his insights with us regarding the practical deployment of deception in the enterprise. As Tushar explains, this is an exciting aspect of cyber security, because it is one of the success stories for researchers developing techniques that find application in live production settings. Here is what we learned from Tushar:

EA: Why do you suppose that deception has not been more extensively deployed in every network to date?
TK: You are correct in the introduction to observe that deception coverage can certainly increase across enterprise and service provider networks. Many teams are just now discovering the power of using deception to detect and prevent attacks. But the method is now officially mainstream, with most enterprise security teams adopting some form of deception. It is even referenced in the NIST 800-53 controls SC-30 Concealment and Misdirection.

EA: When a customer decides to utilize dynamic deception, is the motivation to catch hackers in the act, or to somehow prevent attacks from occurring?
TK: The typical motivation among our Attivo Networks customers is both detection and prevention, where the difference is simply how quickly in the attack chain the deception kicks in. If an intruder is detected early enough in the malicious activity, then the attendant response activities will be performed on indictors, which provides time to avoid consequences. In addition, if the deception is properly designed, then it can divert malicious energy and exploits toward bogus assets rather than real assets. This is a powerful preventive solution.

EA: Do IT teams experience operational challenges in the presence of a deceptive tool deployment?
TK: If the deception is deployed using recommended practices, there should be no collision or operational challenges with IT tools such as scanners. Attivo's solution has the capability of being invisible to these tools, so that we don't trigger them or trigger on them. Additionally, dynamic deception technology delivers operational efficiency with machine self-learning that automates deployment and the adaptive refresh of deception environments.

EA: Can hackers figure out that deception is in place and somehow evade its reach?
TK: The earliest deception prototypes developed years ago using poorly designed lures and honey pots might have been easy to spot. But deception design has advanced considerably to the point where it is indistinguishable from real assets, services, and data. Authenticity and attractiveness are core to modern deception technology. Attivo deception is designed to not only take advantage of the element of surprise, but to also apply advanced deception techniques that are designed to outmaneuver the anticipating attacker. The effectiveness of this has been validated by pen testers who continually fall prey to deception traps, even when they are aware that it is installed within their environment.

EA: What trends do you see in offensive hacking and do you think the defenders are losing ground in cyber?
TK: The clear trend is that hackers are getting better and the barriers to entry lower. This is validated by the increasing sophistication and frequency of breaches in 2017. The impact is significant and the statement that prevention alone is not sufficient, cannot be emphasized enough. Our team at Attivo Network is focused on turning the table on the attackers with adaptive deception that is so authentic that it becomes impossible for attackers to resist and reverses the growing momentum of information security attacks.

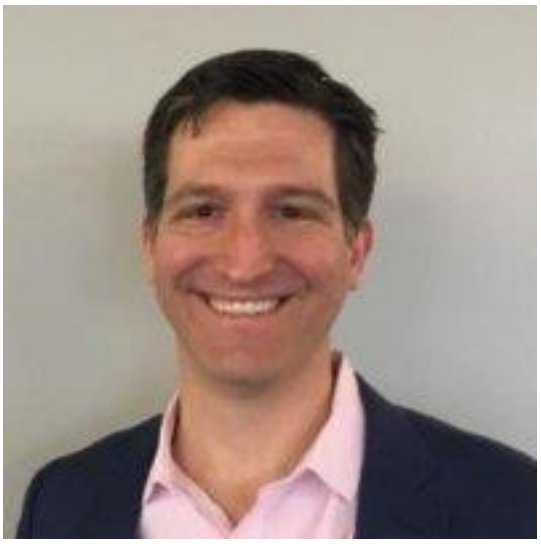

Reinventing Security Awareness: Assessing Risk and Changing Behavior

How predictive analytics and effective training tackle the human error problem

Michael Madon, CEO of Ataata

Employees represent the greatest cyber risk to their companies. Even the most sophisticated technology in the world won't prevent regular people from clicking on bad links, creating hackable passwords, sharing sensitive information with the wrong people or practicing poor office hygiene. We know that most security breaches can be tracked back to human error. So, what can security professionals due to combat the problem? According to Michael Madon, CEO of Ataata, better security awareness training is only part of the solution. Madon's startup gives companies the ability to assess and predict risk among individual employees and across organizations at large. Smarter methodology and an impressive platform demonstrate that Ataata is approaching security awareness in a whole new way. We sat down with Michael Madon to discuss the importance of assessing risk and how Ataata's training provides the antidote to human error.

EA: Michael, how can Ataata help companies analyze employee risk?
MM: Security professionals know how difficult it is to separate the noise from true warning signs. If logging in to the company's network after midnight triggers an alert, should a CISO be concerned every time that bell sounds? Or every time someone prints a document after 10PM? Or any other action that might get flagged? False positives abound, and it's difficult to clear the clutter surrounding an employee's digital footprint. It's critical that we understand why one person's behavior might be more problematic than another's. Ataata can help security professionals identify immediate issues by assessing employee risk and analyzing notable deviations to predict future problem areas.

EA: Explain the connection between employee risk and Ataata's security awareness training.
MM: We know our training platform is changing the way employees think about security and making a real difference for our clients. But at the end of the day, it all comes down to

data. Our proprietary algorithm analyzes sentiment, engagement, and knowledge. And every time a user interacts with the platform, it tells us something valuable. Ataata's platform continuously collects information about individual employees, specific organizations, and industries at large. Are certain employees completing their learning modules right away? Who has missed three knowledge questions in a row? Which employees seem susceptible to specific types of phishing tests? Do employees in one department have different attitudes about security than another? Are certain industries better about protecting information than others? Data from one event won't tell us much. But data collected from millions of events over time can reveal incredibly valuable information for the companies we serve.

EA: You said Ataata's security awareness training is making a difference. What makes your approach more effective than other solutions in the marketplace?
MM: A lot of security awareness programs miss the mark. The challenge is finding a way to engage people who've been conditioned to tune out training completely. You can't expect to change employee behavior if you don't give them a reason to care about security. That starts with better content and a seamless user experience. Our learning modules are written and produced by veterans of the TV industry. The result is content that looks and feels like a typical workplace comedy. Once we capture their attention, then we can show employees how their security decisions can impact the company and their own lives.

EA: Take us through the user experience.
MM: Ataata's model is software-as-a-service (SAAS). Each month, employees receive a new learning module. They watch a video, answer a question designed to reinforce the key takeaway, then see how they performed compared to their colleagues. Companies can also customize their training by including their own material relevant to that month's topic. It's quick and seamless. The whole process typically takes three minutes to complete. And since the platform is universal, employees can complete their training wherever they are from any device.

EA: How do you find the right balance between entertaining your users and getting them to change their behavior?
MM: Our top priority is keeping our audience engaged. The challenge is presenting everything in a way that any employee can relate to and understand. Our creative team consults with industry experts to make sure we're covering the most important topics and conveying the right messages. Our goal is to get people to pay attention and understand why their actions matter. Most security breaches involve human error. If we can get employees to take a step back, breathe and think about what they're doing, we can help companies reduce those breaches and ultimately save time and money. And we know it's working. After three months of training, 83% of Ataata users indicate that they've changed something in their daily lives to make them more secure.

EA: Talk about the Ataata dashboard and how security professionals can monitor their training.
MM: At Ataata, we put as much thought into our analytics as we do into our content. We know that awareness training only has value if you can measure its effectiveness. Our

dashboard gives CISOs the ability to see who's watching their videos and who's answering their knowledge questions correctly. We also assess employee attitude about security and track how that changes over time. And we show companies how their workforce is performing compared to people in their industry specifically and compared to all industries in general. Much of the information available on our dashboard plays a critical role in how we assess employee and company risk. That's by design. Training gives us better insight into risk. And understanding risk helps us develop more effective training.

Mitigating DDOS Attacks in the Modern Enterprise

Significant recent advances in botnet-originated DDOS attacks require new mitigation solutions for the enterprise.

Darren Anstee, CTO of Arbor Networks, the Security Division of NETSCOUT

As botnets have continued to support newer and more lethal forms of DDOS attacks, enterprise security teams have scrambled to react. The traditional approach of redirecting traffic to scrubbing complexes continues to be the primary network architectural means for protection, but the process continues to become more complex. IoT-based botnets, for example, exhibit different characteristics than PC or server-based DDOS origination. As a result, enterprise teams must partner with the most experienced security companies to developed effective solutions. Darren Anstee, CTO of Arbor Networks, sat down with us to discuss recent trends in DDOS being observed and dealt with by the company.

EA: How have DDOS attacks evolved over the past few years?
DA: There have been many key changes in the DDoS threat over the past few years. First, we've seen peak volumetric attacks sizes climb alarmingly, up 60% year-on-year to 800 Gbps in 2016. But it's the frequency of larger attacks that has really grown most alarmingly. Just a few years ago, Arbor's ATLAS systems only monitored a handful of attacks annually above 100 Gbps. Last year, however, we tracked more than 500, which was greater than double the number we monitored in 2015. And this was, in turn, *double* the number we saw in 2014. Second, we've seen more sophisticated multi-vector attacks proliferate. Multi-vector attacks involve an attacker launching multiple attack vectors at the same target, at the same time, possibly targeting different aspects of their infrastructure and services. Arbor's World-Wide Infrastructure Security Report (WISR) shows that the number of ISPs seeing these more sophisticated attacks has increased from roughly a third in 2014, to a half in 2015, and now two-thirds in 2016. Both these changes have been driven by the weaponization of DDoS. Monetized DDoS services remove the need for any technical knowledge when generating large volumetric or sophisticated multi-vector attacks. Multi-vector attacks previously required an attacker with significant resources and skills. But now anyone can launch them, and this has changed the game at a time when many more

end-user organizations are becomingly increasing reliant of Internet connectivity to access cloud, SaaS, and other modern services.

EA: Do you see DDOS attacks taking more advantage of cloud infrastructure either as targets or botnet hosts?
DA: Data centers and hosting providers have been magnets for DDoS attacks for many years. Every year in the WISR, we see a higher proportion of respondents tell us that their data centers have been attacked, that the frequencies of attacks are rising, and that they are seeing more attacks that saturate their Internet connectivity. In 2016, nearly three quarters of data center operators told us they had seen DDoS attacks that impacted their service delivery. Larger cloud operators are already well defended, so it is the smaller operators that really need to ensure that they have the right services and solutions in place. When it comes to using cloud infrastructure to launch attacks, again the larger cloud operators have mechanisms in place now to detect and prevent this. Again though, we have seen smaller cloud operators and hosting provides having their infrastructure abused in this way.

EA: What's been the effect of IoT on DDOS attacks?
DA: On the non-technical side of things IoT, the attacks against Krebs, Dyn, and others last year made DDoS a risk that needed immediate re-assessment for a range of organizations. Many enterprises became acutely aware of the threat and put projects in place to identify and manage their risk. Many ISPs re-assessed their DDoS monitoring and mitigation capabilities given the changed threat. More technically, the most obvious problem is the level of capability that is available to attackers through the compromise of IoT devices. There are millions of devices out there that could be leveraged, with the very large volumetric attacks seen thus far only utilizing a fraction of that. Application-layer attacks have become the most prevalent form of attack from IoT, targeting DNS or HTTP/HTTPS services, and the weaponization of IoT botnets has made access to these cheap resources simple. The less obvious problems, as you alluded to earlier, are those around how we defend ourselves. In most ISPs, the DDoS monitoring and mitigation capability is designed to look for traffic coming in from the Internet, so that DDoS attacks can be identified and stopped before there is any service impact. IoT changes this. It is now possible for the devices connected to an ISPs network to generate enough outbound or cross bound traffic to cause service issues, and this has driven some ISPs to alter the way in which their monitoring and mitigation infrastructure is deployed. Enterprises aren't safe from this problem either, although the issue is much less well-known. We have seen Windows malware designed to scan enterprise internal networks for IoT devices, so that they can be compromised. This could lead to organizations having compromised IoT devices inside of their perimeter controls. The defenses they have arrayed to deal with the DDoS problem would be in the wrong place.

EA: Do you think DDOS volume will ever reach full ISP peering capacity, say, in the United States – or is this unlikely?
DA: It is almost certainly possible for an attack of large enough magnitude to be generated now to hit that level. It is unlikely that this will happen though, at least currently. We must remember that the infrastructure used to generated DDoS attacks is monetized, as in DDoS services for hire. It is in the attacker's interest to ensure they can re-use their infrastructure

as much as possible. Launching DDoS attacks against individual targets, that are large enough or sophisticated enough to achieve their goal without impacting ISPs is the best option for the attacker. Launching attacks that cause broader problems across the Internet, or within an ISP, attract significant attention both from the operational security teams within the ISP community and law enforcement. Obviously, the attacker doesn't want this kind of attention. Where this doesn't hold so true involves ideological hacktivism or nation-state activity. If we do see attacks of this nature in the future, it is likely, in my view, that these motivations will be behind them.

EA: Should smaller companies be considering DDOS solutions?
DA: Yes, although they need to be packaged differently. Smaller organizations are adopting cloud and SaaS even more quickly than larger organizations, partly because they have identified that they have the same risks around data-theft, but no resources to manage that risk internally. These technology shifts mean that smaller organizations are becoming increasingly dependent on connectivity for access to mission critical systems. If there is no connectivity, in a lot of cases, they can't carry out day-2-day business operations. Smaller organizations need to ensure that their cloud and SaaS providers have sufficient DDoS protection in place for their data-centers, and they need to ensure that their own connectivity is protected via their ISP or a cloud anti-DDoS service. These services need to be packaged such that they are provide good visibility of the value they deliver, but abstract away the technical complexities of dealing with attacks.

Mobile Device Security with Emphasis on Apps

Holistic solutions to protecting enterprise mobility must never forget that apps remain a central risk component.

Domingo Guerra, President of Appthority

Thirty years ago, computer security was mostly about making sure you never loaded a floppy disk into your computer, unless you knew exactly where it came from. Fast forward to today, and we download apps everyday onto our mobile – and none of us really know where the software came from. Granted, Apple and Google perform valiant tasks to try to identify evidence of privacy violations or other unwanted functionality. But the fact is that app risk remains a significant issue, even after passing through the filters of the app stores. Domingo Guerra, President of Appthority, spent some time helping us understand the company's unique approach to app risk in the context of their holistic approach to protecting mobile devices and systems from security vulnerabilities.

EA: Let's start with mobile risk: What should enterprise teams be concerned with today?
DG: Enterprise mobility has been in constant flux over the last decade. In that short timeframe, we saw a transition from Blackberry dominance, to secure email and containers on iOS and Android, to the rise of MDM, their transition to EMM and now their transition again to UEM. Early on, some IT and security teams believed that these device management solutions would be enough to secure their mobile environments, but the considerable rise of mobile risk quickly led them to realize that while MDM/EMM/UEM are great management tools, they are not security solutions. This gave rise to Mobile Application Reputation Services and Mobile Threat Defense solutions. This makes perfect sense, because users began to leverage their mobile devices for more than just email. And as the amount of sensitive corporate data and systems that could be accessed form mobile devices increased, so did the need to protect them. However, I think that there remains a bit of complacency among some teams, perhaps with the expectation that a massive and serious attack must occur locally before proper risk mitigation is put in place for their mobile ecosystem. Malware, data leaks, and privacy loss are happening daily via mobile apps and connections so complacency is not a strategy we would advise.

EA: Do you see apps as a primary threat vector for more insidious attacks than what we've seen to date?
DG: Our mission at Appthority is to prevent this from occurring, but we do see the possibility of more dangerous attacks occurring, simply because citizens, business, and government have become so much more dependent on mobility to function. You are correct to point to apps as a primary attack vector, simply because there are so many of them, from so many different sources, with such varying degrees of risk. And app risks are not just the future, but the reality. Apps are the way we communicate and they are constantly accessing and sharing a lot of data. Much of this data is valuable making app-related data a target for hackers. Also, the rise of third party developers and the advertising economy means security is least accounted for in apps, compared to devices or networks. Today, most attacks already involve an app install, whether the user is tricked into side-loading an app from a third-party app store that was never reviewed by Apple or Google or they unknowingly download a malicious or risky app from an official app store which evaded detection. At an OS level, apps have more access and permissions than mobile browsers, so native apps usually do more harm than HTML5 apps or websites. While most mobile attacks to date have targeted the individual (the employee) for personal data, the fact that there are now more mobile devices than laptops and desktops worldwide, and the fact that we are increasingly using these devices for work, explain why we are seeing a rise in targeted enterprise mobile attacks; mobile fleets are often less protected and traditional enterprise infrastructure.

EA: Appthority has really extended its solution to a more holistic approach to mobile security. Tell us about the functionality you've introduced.
DG: With our original successes in the marketplace helping enterprise security teams accurately measure and manage mobile app risk, we realized that our solution could easily extend to a more holistic risk management solution for mobile devices, networks, and apps. And that's exactly what we've done with our comprehensive enterprise Mobile Threat Protection solution. We now provide an enhanced level of enterprise mobile security that deals with mobile risk from top to bottom, and allows us to lead from our strength — mobile app security – and protect our customers from all mobile threat vectors. Further, it allows us to not just reactively address threats with active on-device protection, but to implement proactive security measures of mobile app risk management. For example, if we know an app will be widely used by employees to access corporate data or systems, we can ensure the app has properly leveraged certificate pinning, so that it will not be susceptible to network based attacks like MITM. In other words, we don't have to wait to detect a breach in progress, we can prevent it entirely.

EA: The relationship between MDM and mobile security has always been an uncomfortable seam for many enterprise security teams. Are there ways to smoothen this interaction?
DG: One of the main functional requirements used in the design of our solution was full connectivity with available security tools including EMM, MDM, SIEM, and other tools. No security team exists in a vacuum and no one has a greenfield, so we knew from the start that interoperability and sharing capability to other mobility tools was essential. Further, we see that mobile security is no longer seen as its own silo, but as part of the overall enterprise security strategy. Thus, Mobile Threat Defense solutions need to be able to

inform the right teams *and* systems about active mobile threats. There is also a hidden benefit to leveraging MTD with EMM. As mobile deployments grow, managing and enforcing traditional EMM policies across the organization becomes a difficult manual task. However, by adding MTD, IT and security teams can automate not just detection, but the remediation of mobile threats while simultaneously providing on-device notification and education to employees. This helps lower the IT burden of managing a mobile fleet, and make employees part of the solution, not just a liability. Our goal with our Mobile Threat Protection solution is to add a layer of mobile threat intelligence that enhances protection and informs and automates compliance and remediation.

EA: What do you see as the interdependencies between mobile and cloud security? Can they be treated separately?
DG: Obviously, there are some differences in the types of protections that are embedded in each area, with CASBs and micro-segmentation solutions leading the pack for cloud. But you are right to point out the clear interdependencies between mobility and cloud. For example, cloud security tools have come a long way and can protect mobile traffic while users are on corporate WiFi. But most cloud security solutions have a blind spot for mobile traffic over 4G/LTE. In fact, depending on your perspective, you can think of cloud infrastructure as rounding out the mobile experience – or alternatively, as mobility as a window to cloud. Either way, the security implications of a hack to mobile or a hack to cloud tend to have cross-ramifications. We advise our clients to properly attend to both areas.

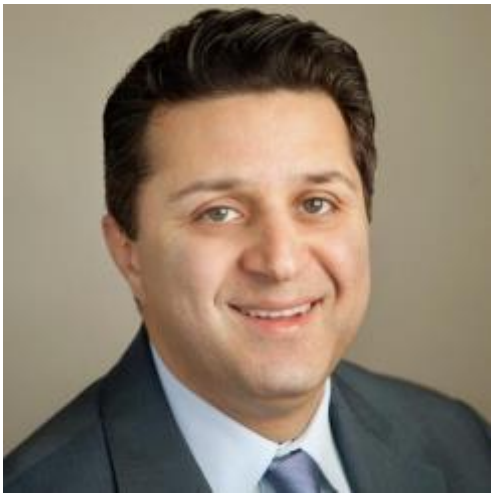

Unified Security for Threat Prevention and Response

Improving the effectiveness of how enterprise security teams detect and mitigate advanced cyber attacks

Barmak Meftah, CEO of AlienVault

Now that virtually every substantive company on the globe has some type of SIEM processing, it is easy to forget that the technology is relatively new in the context of IT infrastructure. Only a decade and a half ago, companies were still trying to determine how to deal with the growing problem of data, telemetry, and security alarms being generated from the local computing environment. More recently, however, InfoSec teams have gotten much better at processing and using collected security information more effectively, not just to detect indicators, but rather to enable threat prevention and to optimize incident response. We had the great opportunity to sit down with Barmak Meftah, CEO of AlienVault to learn more about these trends and obtain his insights into future directions in our industry.

EA: Barmak, let's start with some observations on existing enterprise security: Do you see challenges in the way current CISO teams detect attacks and respond to incidents?
BM: The game is changing today in how InfoSec teams detect attacks and respond to incidents, and the drivers are both offensive and defensive. First, we know that malicious actors have grown more capable in their exploitation of vulnerabilities. Nation-states, for example, have dramatically increased their offensive cyber capability, and this has had an impact on businesses of all sizes and sectors. Second, we also know that enterprise architectures have gravitated to the cloud, especially in the mid-sized enterprise. This shift has changed the nature of cyber security, simply because the underlying computing, network, and application architectures have been virtualized.

EA: What role does automation play in improving both detection and response?
BM: Automation is essential for both detection and response. Automation enables a faster time to discovery and enables IT security teams to respond faster, reducing the mean time to response. The faster a team can identify and respond, the less likely a chance that a threat actor will penetrate and do meaningful harm to the business. Our team is focused on creating unified security management solutions by integrating essential security

capabilities into one platform powered by real-time threat intelligence from AlienVault Labs. This results in a range of capabilities for asset discovery and inventory, vulnerability assessment, intrusion detection, behavioral monitoring, SIEM, and log management that enhance speed to threat detection and reduced time of incident response. Automation is a critical element.

EA: Tell us about the evolution of the SIEM. Do you see the technology evolving to support progression to cloud and mobile?
BM: With today's threat environment, a SIEM is not enough and organizations of all sizes need to evolve from deploying point solutions for interpreting log information into a real-time unified detection and response platform that includes the necessary security controls including asset discovery, vulnerability assessment, intrusion detection, behavioral monitoring, SIEM, log management, and threat intelligence that serves as the primary control point for most security architectures. This primary control point has become even more useful for InfoSec teams as they have moved to cloud, where the associated security information and resources become available on-demand.

EA: Any thoughts on how advanced algorithms can improve threat analysis and response? Does machine learning, for example, play a significant role?
BM: Algorithms for correlating and fusing data, combined with real-time threat intelligence, have improved from simple matching of known attacks to behavioral monitoring techniques that can identify meaningful indicators in large volumes of data. The whole idea is to create actionable guidance for security managers, so the best available techniques are essential. Today, this means understanding the operating environment, identifying changes to that environment, spotting suspicious behaviors on the physical network, the cloud network, on systems, and within applications that can be indicative of threats.

EA: What advice do you have for companies who might have less resources for cyber security?
BM: Security management can be a complex, time-consuming, and expensive undertaking for all organizations, but especially for those with limited security resources, time, and budget. But it doesn't have to be. We founded AlienVault to help organizations of all sizes achieve the same threat detection and incident response capabilities as Fortune 500 companies without the headaches and hassles of deploying, integrating and managing multiple products.

Infrastructure Controls to Secure Enterprise Email

Advanced tools exist to improve the security and authenticity of email senders across the Internet

Ravi Khatod, CEO of Agari

Even after decades of development, billions spent in capital, and an impressive array of security defenses that protect the technological assets of their enterprise, many CISOs remain stuck fighting a defensive battle against digital deception that targets humans. Today, 95% of security breaches start with a deceptive email. The damage from just one of these attacks can be catastrophic both in their immediate financial consequences and in the long-term erosion of trust in business. But the long-term impact on productivity, brand, reputation, and trust in business are even worse. It's not just your identity, brand and employees that are being targeted. The very fabric of trust in digital business is at risk. But it doesn't have to be this way. What if you could turn the tables on cyber criminals and shift from reactively defending against threats to proactively securing your business? What if you could spare your people from constantly being at risk and having to waste time and energy deciding what they can and can't trust, and they could just focus on their jobs? And what if instead of continually trying to keep up with the latest forms of attack, you could put an AI-driven system in place that understood the fundamental strategies hackers use and get one step ahead of the next big hack? Ravi Khatod, CEO of Agari sat down with us recently to share his views on these important issues.

EA: Why do we continue to see so many email-based threats despite all the security controls in place (e.g., secure email gateways, sandboxing)?
RK: Unfortunately, cybercriminals have figured out how to bypass those defenses and target the most vulnerable part of any enterprise: Humans. People and things pretending to be someone or something they're not are slipping past our defenses daily. Spear phishing and other identity deception attacks over email are the current attack vector of choice, accounting for 95% of all security breaches. And this problem is only getting worse as more companies embrace digital business models and create an increasing number of digital attack vectors for cybercriminals.

EA: Who specifically should have the responsibility to protect against this kind of digital deception? Is it service providers? Businesses? Individuals?
RK: Obviously, the email ecosystem purveyors – usually a combination of the IT and security teams – in an organization will have primary responsibility for establishing increased trust and security for sending and receiving email. The mistake is to put the onus on individuals to constantly be looking over their shoulders or interrupting their work streams to try to spot bogus emails. That interferes with productivity, and beyond that, it sends the wrong message to people—that they can't trust digital business. We're not just talking about defending against an attack over email—we're talking about defending against the very fabric of trust in business itself. And that should be the shared responsibility of not only IT but everybody involved in the business.

EA: What solutions do you offer customers for improving trust and security in their email infrastructure?
RK: Agari protects people and businesses against cyber criminals that use false identities to commit fraud, steal information, and undermine trust in digital business. The Agari Email Trust Platform is the industry's only artificial intelligence (AI) driven defense system to protect humans from being deceived by cyberattacks such as phishing, ransomware, and business email compromise. What makes us unique is a combination of our strategic approach and how we leverage AI. Instead of trying to anticipate and block all the unknown ways to attack an email system, we're able to use our AI-driven defense system to model what good, trustworthy digital communications look like and use that to recognize and block fraudulent ones. The result is not only the ability to trust your email, but an ability to trust your digital business as it continues to grow.

EA: When email spoofing is detected through advanced monitoring, what sort of mitigating action can be taken?
RK: When you think about it, the old model of incident response teams and "mitigating action" falls short of real protection. The goal of our solutions is to stop targeted email attacks before they can occur. Think about it. What if you could leverage AI to identify what trusted communications look like and use that as the means by which imposters were recognized and blocked? What if people—and I mean both end users within an organization and the IT and security departments—spent less time cleaning up after security breaches and spent more time on strategic considerations? What if you could just trust your business communications rather than wasting your time on dealing with the lack of trust? It's time for the security industry to turn the tables on cyber criminals and use AI to model what's trustworthy rather than constantly playing catch-up trying to figure out what isn't trustworthy.

EA: Any predictions about the future of the email security threat in the coming years?
RK: Today, cyber criminals behave more like well-run businesses than lone hackers. They use the most cost-efficient form of attack that produces the most return. Email is their communication vector of choice. Email is ubiquitous and unauthenticated, which makes it easy to deceive victims by using false identities. As security controls implemented IP reputation and malware signatures to detect attackers, these criminals pivoted to more targeted types of attacks with no payload and larger payoffs like Business Email

Compromise (BEC), often involving wire fraud, invoice scams, or W-2 scams. This trend will continue its 1000% per year growth because none of the traditional security controls can stop it. In fact, at Agari, we have built with our customers a taxonomy of known attack types and methods that are most frequently used by cybercriminals—a taxonomy that translates not only across email but other forms of digital deception as they evolve. And it's a playbook that helps our customers, and us, stay one step ahead of the bad guys.

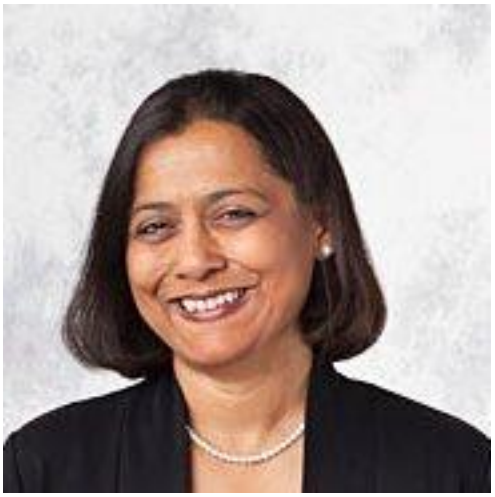

Monica Pal, CEO of 4iQ

Identity Threat Intel in the Deep and Dark Web

Advanced scanning of surface, social, deep and dark Web sources to detect digital risk

Here is a sobering fact: Over three billion passwords were stolen in 2016. Common statistics suggest that most people shuffle through about 2 – 5 passwords to access 25 or so online sites on a regular, on-going basis. If we wanted to make things any easier for hackers, we'd have a hard time doing so. With more than 80% of attacks being initiated using stolen credentials, this is quickly emerging as one of the biggest threats to security on the Internet. To address this problem, advanced techniques have emerged for deep investigation of available intelligence for evidence of compromised identities. The best approaches make use of automated crawling and big data analytics, combined with trained experts who know how to navigate the deepest and darkest parts of the web. Monica Pal, CEO of 4iQ sat down with us recently to explain how her company provides fresh identity threat intel so that companies can alert consumers, customers, executives and employees as soon as stolen passwords and exposed personal information is discovered in the deep and dark web.

EA: Monica, what are the risks to digital identities that individuals face on the Internet today?
MP: Consumers face serious threats associated with exposed credentials. If you are like me, you've created accounts on Internet sites that you don't even remember. On top of that, most people reuse usernames, often their email address, as well as passwords across their online accounts. If they are forced to reset their password, they will rotate through a small set of passwords or simply add an additional number or character. Hackers know that your Hotmail or LinkedIn password is probably the same as for Dropbox and banking, so they use credentials stolen from one account to test and unlock other accounts. Once they take over, say, an email account, they could have access to conversations, chats, contacts, calendar, documents, photos, and more. They can invade your privacy, learn about who you are, determine where you live and what you think, access your calendar, publish conversations and photos, or use information for social engineering. They can spam your

email, access your social contacts, send phishing messages, and infect them with malware and ransomware.

EA: Do businesses face similar risk? And do you see these risks increasing with social media?
MP: Yes, businesses face similar risks. The lines are blurring between personal and business, as well as on-line and off-line. For example, most people no longer switch phones or tablets for business and personal use, and we all use the same passwords for personal and business accounts. So, although businesses continue to invest money protecting IT infrastructure, a hack on a small, unrelated gaming site can leave the door wide open to the enterprise. Social media increases these risks exponentially. The sharing that occurs with family, friends, and business contacts creates a treasure trove for criminals as they figure out who to target and how to attack. Executives and boards are especially susceptible. Criminals study their family, friends, and business associates to launch what are sometimes called 'CEO scams.'

EA: What are the best sources of intelligence about threats to our digital identities?
MP: Once hackers exfiltrate usernames, passwords, and other online account data, they either use it themselves, sometimes over months and years, or give it to brokers who trade amongst friends, which are rings of anonymous personas talking in IRC channels in the dark web. If you follow this trail, the best sources of intelligence are in close communities of the dark web, where you need to know the right personas and have the right reputation. Next come black markets in the dark web where these data sets are sold, followed by a couple hundred other forums and Twitter handles where information on stolen credentials and personal information packages are exposed. Since digital identities are central to our digital lives, we have focused on searching the surface, social, deep, and dark web, looking specifically for stolen, lost, and leaked data that might contain personal information.

EA: You've mentioned a couple of times now, the surface, deep, and dark web – what more can you tell us about them and how does your platform access these sources?
MP: The surface web is the most common and well-known. It is that portion of the web that we use every day, and is indexed by standard search engines. The deep web, in contrast, is bigger and includes content not indexed by search engines. The dark web is smaller and contains content not indexed and not available via standard browsers. You must use special browsers like Tor to anonymously access sites, forums, and IRC channels. In addition, sites in this part of the Internet are transient. That is, they come and go, sometimes are up and sometimes are down. The more coverage and context, the better the intelligence, so our platform scans all parts of the web, including surface, social, deep and dark. Many parts of the dark web cannot automatically be accessed, so our subject matter experts go into these places and manually monitor chatter and collect information. Once data is collected, our system automatically structures or normalizes the data, extracts and disambiguates identities, and stores a hash of the information. Customers who have registered a hash of their digital identity with us are sent an alert as soon as new exposed information on them is found. This allows them to change passwords, adjust privacy settings, reconfigure servers, and limit damage.

EA: Won't passwords soon be a thing of the past? Aren't people going to use two-factor authentication (2FA) and then this problem will be gone?
MP: The problem may have more to do with human nature than technology. For example, two-factor authentication has been available for decades. But it is hasn't been easy or cheap. Even today with mobile devices used as the second factor, it is not easy for businesses to simply move to 2FA. It is a speed bump that could turn off users and negatively impact the bottom line and for many businesses and, given the choice, very few people will turn 2FA on. But even if 2FA was widely adopted, clever criminals can trick users into sending the PIN to the hacker.