the cyberwire

The CyberWire
621 East Pratt Street, suite 300
Baltimore MD 21202–3140
editor@thecyberwire.com
www.thecyberwire.com
@thecyberwire

RSA 2016

February 29 - March 4, 2016 San Francisco, California

RSA Notes—Perspectives on Threat Intelligence

Threat intelligence can be easily conflated with attribution. "Who did that to us" is a natural first reaction in an enterprise that's been hacked.

But while we heard a great deal about threat intelligence at RSA, almost without exception we heard much less about attribution than we have in the past. There's a general conviction that actionable threat intelligence is vital to security, but also a general sense that much of what has passed for threat intelligence hasn't, in the end, turned out to be particularly valuable.

That cyberspace has become an invaluable domain for national intelligence services is beyond question. Collecting in cyberspace is an obvious natural technical evolution of the long-established discipline of signals intelligence. But the sort of threat intelligence we heard discussed at RSA is intelligence developed and applied, for the most part, by businesses, not states, and it's in the first instance intelligence about threats to enterprises. These are mainly but not exclusively the familiar sorts of cyber threats: data exfiltration, denial of service, and so forth.

Taking the intelligence metaphor seriously, it's worth noting that the intelligence cycle classically runs from direction through collection, processing, analysis, dissemination, and feedback. So raw data become intelligence only upon analysis. They're also collected with some purpose in mind—the "direction" phase of the cycle—that is, intelligence should be actionable.

And actionable intelligence is seen by those who provide it as genuine intelligence: collected, processed, analyzed and disseminated with some clear purpose in mind. It is not a mass of logs, chattering alarms, or unanalyzed data. Such data are fatally easy to collect, and in many ways have replaced Clausewitz's famous fog of war *(insufficient information)* with a glare of war that's equally blinding.

Threat intelligence is interesting and valuable insofar as it reveals what an adversary is trying to accomplish and what tactics, techniques, and procedures they're likely to use. Understanding these can usefully inform the organization of defenses. As A.J. Shipley of LookingGlass put it in discussion with us, "Actionable intelligence is something that reduces your risk profile. If it can't reduce your risk profile, then it's not actionable." And if it's not actionable, then it's not worth much.

Our talks with companies exhibiting at RSA suggest that some of the common themes on people's minds include the importance of context and actionability in the development of threat intelligence. Many of the companies specializing in threat intelligence work from unstructured, open-source data with a view to providing either insight into probable adversary goals *(and the tactics, techniques, and procedures they're likely to use in pursuing those goals—what Palo Alto Networks CSO Rick Howard calls the adversary's "playbook")* or into ongoing attacks. The former is useful in establishing defenses and hardening networks, the latter in detecting, mitigating, and recovering from incidents.

Here are notes from our conversations with companies offering varieties of threat intelligence. *(You'll hear interviews with other organizations in the CyberWire's Threat Intelligence Podcast, posted on March 8, 2016.)*

ProtectWise—"a network security DVR."

ProtectWise characterizes itself as a "network security DVR." It offers a software agent installed on a network port, optimized for high throughput packet capture. ProtectWise aims at full packet capture, moderated by policy *(dropping protocols without security relevance)*, then compressing, encrypting, and uploading the data to the cloud, where they're then run against a suite of analytical tools.

The system's aim is detection, visibility and response. ProtectWise distinguishes its offering sharply from the sort of log management one would find in a SIEM, calling it a "wisdom engine." Uploading it to the cloud enables collective correlation, but no data are shared among customers. Rather, the system does packet discovery.

One of its more interesting capabilities is its ability to retrospectively discover zero-days. James J. Treinen, Vice President, Security Research, described five kinds of use cases for the technology.

1. Real time monitoring. Packets stream across all the threat detection capabilities deployed.
2. Retrospective analysis. In the event of a major exploit involving an indefinitely large number of potential victims, an enterprise can return to its data and determine whether or not it was affected.
3. Deep forensics. Since all network traffic is preserved, the records of that traffic can be used when extensive forensic work is necessary.
4. Case management. The system's dashboard provides a perspicuous interface for security and IT teams.
5. Network administration. The system offers a way of doing deep profiling.

The physics of ingesting such large volumes of data were particularly challenging during development of the system. "So is storage," Treinen noted. "You can't simply brute force that much data into storage. And retrieving data is also challenging. Both storage and retrieval required them to develop massive parallelizing techniques.

Having all the packets not only assists in deep forensics, but also helps responders triage incidents. It helps them move from the interesting to the suspicious to the malicious.

This is not, however, simply indiscriminate collection. The system is designed to process the very large volumes of data it collects into a form that human analysts can use, reducing their workload rather than overwhelming them. The strong cues the system yields helps incident responders to pull the packets they need for forensics and remediation. This is an advance over old school reactive cues—strong hints tell teams where to hunt.

Many of ProtectWise's customers come from the entertainment vertical. One of the major customers they can discuss is Netflix, which has become a major partner. Netflix sets the template for the customer base—large enterprises with strong existing security capabilities are a good fit for ProtectWise and its "wisdom engine," which is well-adapted to making mature incident response teams more effective.

Treinen himself has a background in false positive suppression. The company's "wisdom engine" is designed to filter out the glare of too much information. It's intended to be a machine smart enough to look for interesting patterns, and sufficiently well-adapted to analytical needs to enable an enterprise to scale advanced human talent. "It treats every output as an input. The machine needs, in the end, to give an input to a human."

The system avoids evasion by threat actors because it casts a very wide net. It checks signatures *(including through a classical IDS)*, then applies heuristics, machine-learning classifiers, and finally behavioral analysis *(and this behavioral analysis can be analysis of either machine or user behavior)*. Its visualizer is designed to help the noise fall away, using ambient display theory.

Some cases of interest include investigation of xCodeGhost—they found it went much farther back than initially believed. Retrospective discovery of credential threat has been interesting. They've found APT activity early in the establishment of persistence as threat actors stood up their infrastructure. In fact, retrospective discoveries are often the most interesting, because they often reveal the most closely targeted attacks.

"The network never lies," Treinen concluded, and capturing the ProtectWise's vision is to make incident response as effective as possible.

LookingGlass—an Internet-perspective on defense.

A.J. Shipley, VP Product Management at LookingGlass, took us through some of his company's capabilities. LookingGlass has grown through acquisition, most recently acquiring Cyveillance from QinetiQ North America. Shipley noted that this acquisition was strategic, not opportunistic: "Our intention has been to build a trust-centric, next generation security company," bringing together data collection and analytical capabilities with the goal of "giving the human analysts their bandwidth back."

The capabilities they picked up with the Cyveillance acquisition gives LookingGlass the ability to see the information the bad actors are exchanging. With its unstructured open-source collection capability, Cyveillance gives rich, machine-readable threat intelligence to users at various levels of sophistication. LookingGlass started with an Internet-perspective. This is, in effect, terrain and weather. Cyveillance gives the enemy situation. Scout delivers the friendly situation. Analysts who receive the intelligence can see a great deal; subject matter experts presented with it can easily extract relevance, Shipley said.

Originally, in the intelligence disciplines, cyber was a medium in which people launched attacks or exfiltrated data. But now it's seen as a space within which we collect intelligence.

ThreatQuotient—intelligence collaboration for security operations.

Ryan Trost, CTO and Founder of ThreatQuotient, explained that his company emerged from its founders' experience in security operations centers. "We used to have intelligence," he said, "but no good place to collaborate." Analysts would store information in Excel spreadsheets, Word documents, and so forth.

ThreatQuotient created a database to consume and ingest threat intelligence, then push it to sensors and SIEMs. This information is provided, Trost said, in context. You may, for example, get an IP address, and be told that someone suspects the address may be malicious, but beyond that you need to know the role the threat indicator plays. This jumpstarts investigation, making triage of reported incidents much faster. The more structured the information, the easier and the better such triage proceeds, but historically the information has been structured in accordance with a wide range of different standards.

The problem of too much information is a real one. Security teams can suffer paralysis by analysis, often using as many as two hundred or so open source feeds. Teams have to build the capabilities they use. But companies need to engage in some self-reflection: the feeds they employ must be useful. ThreatQuotient, Trost said, helps companies determine which feeds are useful.

Knowing the adversary enables you to build better defenses against people with specific motives who use a specific range of tactics, techniques, or procedures. Nation-state cyber operators tend to be relatively systematic, certainly more systematic than the notoriously opportunistic crime syndicates.

Threat intelligence platforms standardize over time the various black bag tricks analysts have. The platform allows analysts to build a playbook that enables collaboration, with situational awareness being the ultimate goal.

ThreatQuotient's primary purpose is to be an on-premise integration tool. We integrate with sensors and SIEMs, and we lessen the burden on analysts who would otherwise be working manually. Trost concluded with this account of actionable intelligence: "Actionable intelligence arises from additional context that narrows where the analyst needs to look."

Level 3 Communications—mapping threats, but doing so with a purpose.

Chris Richter, Senior Vice President of Global Security Services at Level 3 showed us an extremely elegant threat map of a kind that's become familiar to many of us in the industry over recent years. *(And this one really was elegant—it distinguished, for example, among various kinds of traffic analysts might find interesting. Attack traffic was colored differently from probes, to take one distinction often glossed over.)*

But a display, however compelling, doesn't necessarily or even typically provide actionable intelligence. Richter was very clear on this point: "The map isn't the product—it's marketing." The ThreatMap is a highly filtered result of netflow events. What Level 3 delivers for customers is intelligence that serves very specific security needs, principally DDoS mitigation. In their view, data that help mitigate threats or active attacks are actionable. This isn't *("merely")* strategic intelligence, but rather tactical intelligence, collected, processed, and presented in a way designed to tell you when you're being phished, probed, or exfiltrated. Their services also identify shadow IT, open ports, firewalls, and so on, providing useful insight into the friendly situation—aspects of a network's own security posture.

Level 3 can identify botnets and command-and-control servers, an obviously useful intelligence product because it can be used to take down threat actors. *(They contributed to the recent takedown of Poseidon, for example, and have been prominently involved in tracking the Angler exploit kit.)*

Level 3 operates internationally. They maintain scrubbing centers in Sao Paolo, San Francisco, Los Angeles, Chicago, Washington, Dallas, London, Frankfurt, and Amsterdam, and have plans to establish centers in Asia. The threat intelligence they develop refines rule sets at the scrubbing centers.

As the company looks toward a forthcoming rollout of other security services, they note that these will not be additional threat intelligence products. Rather, they'll be services that benefit from the threat intelligence Level 3 is already set to provide.

Recorded Future—unstructured data meet natural language processing.

When we asked Levi Gundert, Vice President, Information Security Strategy, at Recorded Future to define threat intelligence, he offered this: "It's the act of formulating an analysis based on the identification, collection, and enrichment of relevant information."

Different consumers of intelligence bring various perspectives to it. Businesses, though, in general care about decreasing operational risk and establishing a competitive advantage, Gundert explained. Threat intelligence is no panacea. Cyber security is being organized against frameworks that represent the accumulation of twenty years of threat knowledge, and security will increasingly be focused on compliance, audit, and policy. "Now," Gundert said, "we need to take the next step—identifying what attackers are up to in near-real-time."

Recorded Future takes a programmatic approach to delivering threat intelligence. They take massive amounts of unstructured text, to which they apply proprietary natural language processing. This approach scales in ways that alternatives do not. They also integrate with incident response.

There's value in general attribution, but largely because "motivation informs methodology." Understanding what the adversary is after enables an enterprise to shape its defenses accordingly.

Coming tomorrow—perspectives on emerging technologies.

We'll return to some of these themes tomorrow, but in the context of emerging technologies, and in particular a consideration of the ways in which innovative vendors can integrate their products with larger platforms and managed security service providers."

RSA Recap—Notes on Security Technology Trends

From the companies we spoke with at RSA, several trends appear to be driving technology development in the cyber security space.

Above all else is the continued, longstanding push to automate as much of security as possible. This isn't seen as movement toward replacing human analysts or watchstanders—we heard few hints of a push for such replacement, since human talent seems by tacit consensus to be effectively irreplaceable at some level. Rather companies are interested in offering approaches that enable human talent to raise its game. Hence we saw repeated emphasis on solutions that reduced the need to review logs and watch alerts, and that promised to free human operators to look at the big picture and perform the triage necessary to effective, timely incident response.

Machine-learning approaches to anomaly detection seem to be a popular option. These are seen as cutting through noise with relatively low loss of signal. And the ability to ingest and process very large amounts of data was featured by many of the experts we spoke with. Those data are increasingly accepted in unstructured form.

Finally, scalable, comprehensive security solutions are increasingly seen as vital. This trend has a few interesting corollaries. It offers a space for big integrators to offer managed services that cut through another form of noise—the high volume and rate of introduction the market in security offerings sees. Comprehensive managed security services are also scalable, and enable small and mid-sized enterprises to enjoy the security resources formerly seen only in larger, well-resourced organizations—dedicated security staff, SOCs, even IT teams—and to do so in an affordable fashion. The trend also strongly suggests that innovators with new products would do well to develop them into offerings that could easily integrate with large comprehensive solutions.

We spoke with AT&T and Verizon, two of the big integrators, and here's what they had to say.

AT&T—"Sure, we cover paragraph one, but we really try to address the entire op-order."

Jason Porter, AT&T's Vice President of Security Solutions, has a military background, and so we posed a question about threat intelligence to him in military terms. Do you, we asked, see yourself as covering everything the intelligence staff is traditionally responsible for? Specifically, in the classic US Army field order, intelligence is the subject of paragraph one, "Situation." And his answer was, "Yes, absolutely. But we also try to cover the entire op-order."

AT&T provides managed security for businesses. It runs eight SOCs globally, and from those SOCs it leverages data scientists and data engineers. They're finding, Porter said, that customers increasingly want AT&T to take over their SOCs. Their threat intelligence platform is layered throughout, and they tailor the customer's defensive posture to the situation.

This necessarily involves the ability to handle and analyze very large volumes of data—some 107 petabytes daily. They install agents on devices for sensing, build secure connectivity, and offer a full suite of data-loss prevention, intrusion detection, and firewall products. All system inputs feed their big data platform, and they use machine learning as they watch for anomalous behavior.

There's too much dependence on the human analyst as opposed to letting the machines do as much work as possible. AT&T uses machine learning to identify threats with a much higher

degree of precision, and then enables the analysts to validate the threats the machines flag and kick-off the response. "We're now, when confident of a signature, even beginning to let the machines themselves act. There's no longer any need to have analysts review logs. Their time and attention are better spent looking at specific threats."

Customers, Porter noted, have been inundated with tools. They can spend millions on systems, take years to roll them out, and by that time the systems are deployed, they're already out-of-date. AT&T's approach is to build a platform for end-to-end solutions so that they can bring in and integrate the best technology rapidly.

He described three of the many use cases for his company's security solution. It can, through its big-data behavioral analytics, recognize anomalous user behavior indicative of an insider threat. It can recognize data exfiltration—when you start to see traffic from a server interacting with an unusual IP address, you've received warning that data are probably being siphoned out of your system. And the solution can offer early detection of virtual botnet formation, when it sees a robot talking directly to another robot as opposed to a controller.

Like most others we spoke to, Porter believes that the value of attribution lies mainly in what it reveals about an adversary's goals, and the approach the adversary is likely to take to achieve them. He also sees some other value in identifying and interacting with a specific adversary: "Virtualization enables us to use distributed honeypots. This actually helps us shape the behavior of an adversary, so we can better understand what they're after and what they're doing."

Porter sees securing the Internet-of-things as offering some unsolved challenges. "IT is meeting OT in the IoT," and this presents us with security issues. People are now connecting older, formerly unconnected legacy systems, which is one reason for AT&T's alliance with Bayshore Networks.

In sum, AT&T seeks to detect and respond to threats affecting its customers. "We offer holistic solutions. We don't just play at the threat level, or at the endpoint." While only the highest end customers have data scientists, SOCs, and so forth, the holistic solution AT&T offers seeks to bring in best-of-breed technologies to all customers, and to do so not in a custom delivery, but through a scalable, tailorable model.

Verizon—"Finding the threat that's sticking out, and really pulling on it."

The other major integrator we spoke with was Verizon. Vincent M. Lee, Director, Product Management, Verizon Enterprise Solutions, took us through his company's integrated security service. They find that their customers want transparency with respect to their environments, and sound situational awareness.

Verizon approaches network security with statistical and entropy analysis, "Finding the threat that's sticking out, and really pulling on it." Their analytical support enables them, Lee said, to help their customers cut through the tremendous volumes of data, much of it mere noise. They're swapping out a homogeneous SIEM for Splunk, which, Lee said, has become the de facto standardized data stream, with almost all vendors now building toward working with Splunk.

Until now, SOCs have been customized and operated by a dedicated staff, but Verizon's customers are now using its managed platform to scale, thus saving costs. Regular SOC personnel look at detailed events; the SOC manager looks at triage, and the big picture.

Like others with whom we spoke, Lee sees attribution as valuable only insofar as it enables an enterprise to discern intentions, and tactics, techniques, and procedures. "We don't offer threat data sharing," he said, "but we do leverage threat patterns. The Threat Library Team builds out patterns. In the SOC itself, we tweak these patterns—rules—to customer-specific requirements."

Since this is an integrated security service, Lee said they're able to tell their customers, "Don't worry about vendors; we'll find and deploy the products." It's very difficult for everyone to keep up with a rapidly changing, dynamic marketplace, and they seek to remove that load from their customers.

One aspect of Verizon's offering is distributed denial-of-service *(DDoS)* protection. "DDoS protection is no longer carrier-dependent, but has become carrier agnostic. It's not that easy to defend against DDoS—only four or five players offer defense against volumetric DDoS attacks. *(Verizon is one of them.)* Also, volumetric DDoS defense should be available on demand. There's no reason it should be always on." Their hybrid DDoS protection combines on-site always-on capability with failover protection on demand. This kind of protection is now being tied to Verizon's managed security service packages. *(This tie-in is perhaps unique, Lee thinks, and certainly unusual.)*

Cylance—breaking the cyber risk cycle and lowering control friction.

Malcolm Harkins, Global Chief Information Security Officer at Cylance, said that the community is tired of reactive approaches to security. It's also burdened by alert fatigue.

"Our intention is to operate with automated prevention. We want to reduce control friction," he said. Controls are a drag on an enterprise. Enterprises want solutions with as little control friction as possible. If you've got too much friction, notoriously you'll drive people to work around the very controls put in place to protect them. So to lower risk and lower costs, you should try except in exceptional cases to reduce control friction.

Control friction can manifest itself in a variety of ways. It can drive the creation of shadow IT as personnel work around security controls that impede mission accomplishment. And control friction can itself degrade computing functionality.

Sometimes, of course, a degree of friction is desirable, as may be seen in the use of roundabouts for traffic control. They slow traffic down, reduce the severity of accidents, and in general push drivers toward more attentiveness. "But this isn't the sort of positive friction we see in enterprise IT."

Cylance, Harkins said, is different from other companies in that it operates at the kernel-level. It inspects files in milliseconds, looks at feature sets, and then estimates the probability that some file is malicious. This judgment is made on the basis of data science. Cylance detonates suspect files in a cloud, which further reduces the cost of control. It then contextualizes to learn about threat actors and agents.

The goal is to protect the enterprise in ways that enable a business to achieve its mission.

HEAT Software—endpoint security management

HEAT, which has grown through acquisition, is focused on endpoint security management. We spoke with Russ B. Ernst, HEAT's Senior Director, Product Management.

Their unified endpoint management includes client management, enterprise mobility management, and endpoint protection. It's a standardized solution optimized for the mid-market. That market is looking for enterprise-level configurability, but it's under-resourced. A typical mid-market customer might have between 500 and 25,000 managed devices. It usually has no SOC and no full IT operations team, often no full help desk. The endpoint protection suite has its heritage in patch management. They also have a strong OEM program, and a partnership with Qualys. Once deployed, they do ongoing, dynamic monitoring.

What's your ideal partner? "We have the endpoint nailed. Someone in the network space would be an ideal partner."

What unsolved challenges do you see? We have to be right all the time. The adversary only has to be right once. And security can't hinder productivity, as it sometimes has a tendency to do.

What do you want people to know about HEAT? "We're unifying cloud services and endpoint management through workflow management. We focus on the user, empowering them without unnecessarily escalating their privileges."

Zimperium—protection against mobile attacks.

Zuk Avraham, Zimperium's founder, described his company's approach to protecting enterprises against mobile attacks. "Most mobile security products tend to focus on apps," he said. For its part Zimperium seeks to address the device *(Stagefright is an example of a device vulnerability, one exploitable by nation-states)*, the network *(scans, for example, while not necessarily malicious, could indicate the arrival of an APT)* and applications.

Threats in applications divide into the known bad, the unknown bad, and, worst of all, download-and-execute threats. A download-and-execute threat displays no initial bad behavior, thus bypassing many protections.

Zimperium offers a map that displays detected attacks on networks. It uses Zimperium apps installed by users on smartphones, with the phones themselves serving as travelling sensors. They don't sniff network packets *(because they don't establish root access on the phones)*. Instead, they use a zero-packet approach that instead detects derivative impacts of attacks on devices.

The typical Zimperium customer is an enterprise. Its users install the Zimperium app because it warns them when they're encountering a threat. In a common use case, a CISO might receive a report of a threat. This warning would enable the enterprise to disconnect a compromised phone from corporate networks. Or a bank might restrict transactions from compromised devices. Other customers are those most likely to be targets—high net-worth individuals, VIPs, celebrities, and so forth. Detection is based on device behavior.

Cyphort—"We go for as much automated heavy lifting as possible."

Fengmin Gong, Cyphort's Co-Founder and Chief Strategy Officer, talked about his company's approach to advanced persistent threat protection. "When we build a product, we know how to build it from market experience. Our new product is based on the new threat, which can be anywhere, and can occur at any stage." They aim for 100% visibility, deploying an omnipresent collector on a network, feeding a centralized analytical facility. "We try for as much automated heavy lifting as possible, correlating as much as possible." The analysis triggers an action, like blocking. The product itself doesn't block, but rather supports an open API to leverage client tools.

Like other companies offering advanced threat detection, Cyphort uses anomaly detection with machine learning training. They train a model to baseline software normal behavior. They've recently introduced coverage of lateral movement.

Quick detection is vital because, he noted, "The most important thing is immediate containment *(isolate the infected box)*, intermediate steps *(go to check other endpoints on the network)*, and then conduct cleanup, remediation, education and upgrades *(this last being 'hardening')*." Because it's essential to prioritize responses, Cyphort provides a risk score for the incidents it detects.

"Simply saying, 'I feel something's wrong' without saying what's wrong is unhelpful." Cyphort's advanced threat detection solution is designed, Gong says, to do the "automated heavy lifting" that can make a SOC or a SIEM effective.

"We're building the basis of a new understanding," he said. "We want to build a practical tool. Too many people are still not seeing 100% coverage as the beginning point. You achieve visibility through trained machine learning and active detection, tailored to a specific environment."

A look back at yesterday's discussions of threat intelligence.

It's worth noting that similar themes appeared in yesterday's discussions of threat intelligence. Recorded Future's Levi Gundert stressed his company's programmatic approach, and its ability

to apply natural language processing to unstructured data, as a way of providing a security solution that scales. ProtectWise's James Treinen emphasized the potential intelligent, automated big data processing had for enhancing security. And ThreatQuotient's Ryan Trost drew lessons from long experience with SOCs on the importance of being able to handle and prioritize high-volume data feeds.

Coming tomorrow—notes on trade and investment.

In tomorrow's issue we'll take a look at some international approaches, and offer some thoughts about trade and investment.

RSA Notes—Trade and Investment

We had occasion at RSA to speak with representatives of several international firms and government trade missions. Some of the firms we've discussed in earlier posts; we offer a summary of our conversations with three others in today's post.

The United Kingdom was heavily represented at RSA, and we spoke at length with Andrew Williams, their cyber envoy to the United States. You can listen to that discussion in today's special CyberWire podcast. But it's worth noting a few patterns in the UK's very active presence in the cyber security market. The government has taken an active role in the incubation and promotion of cyber startups. There's an obvious alpha customer in the UK, and British universities are also making a substantial contribution to research and development. In any case, listen to the whole thing.

We had an opportunity to visit the German pavilion as well, where we spoke with representatives of that country's Internet industry association, Verband der Internetwirtschaft e.V.. German firms exhibit a strong commitment to international business and a sophisticated understanding of the agreements, policies, and regimes that shape it.

The well-established *(if dismal)* principle that living in a bad neighborhood tends to produce innovative security products and technologies was borne out by what we learned in conversations with representatives from the Republic of Korea and Israel. South Korean companies are fueled by the necessity of coping with essentially continuous cyber mischief from their northern neighbor on the peninsula. *(We've been following developments there this week, as South Korean intelligence services outline recent cyber espionage campaigns mounted by the DPRK, and as the US and the Republic of Korea move to make an already tight cooperative relationship even closer.)*

We spoke with Taeil Cho of the Korea Trade Investment Protection Agency *(KOTRA)*; he stressed the country's openness not only to exporting advanced technology, but also to partnering with international businesses. We note that the US Department of Commerce will be sending a trade mission to South Korea this May; the mission will also visit Japan and Taiwan.

Finally, we spoke with several companies that operate in Israel. Here's what they had to tell us about their technologies.

SECDO—Ending SOC pain by telling the story in context.

We spoke with Shai Morag, SECDO's CEO. His fifteen years in cyber security started in Unite 8200. A serial entrepreneur, he founded SECDO in 2014. The company currently has about twenty-five employees. The Check Point chair led funding of the company.

SECDO is interested in "ending SOC pain." Many solutions are offered to SOCs, and SOCs necessarily deal with many alerts. The way to end SOC pain is to, effectively, multiply the tier-1 analysts. Since everything comes to the endpoints, the solution is to have full visibility, and you can achieve that by having collectors on the endpoints, and taking alerts as leads. An analyst needs full understanding of the story, and SECDO develops the full context to score the threat.

Their approach, said Morag, uncovers threats that would otherwise go unnoticed. You want a system that builds the story for you—you don't want to have to query everything. "We shorten the dwell time between detection and response." And above all, Morag emphasized, it's important to "give context."

In the future, SECDO intends to offer remediation services. Its "Ice Block" is a patent-pending technology that remediates by freezing malware in memory, quarantining the file, and reverting the changes in the registry.

Currently SECDO aims for customers who have a SOC, and they partner with managed security service providers. "We follow the SOC."

SCADA Fence—Securing the production environment.

Yoni Shohet, co-founder and CEO of SCADA Fence, described his company's approach to industrial control system security. They focus on production environments, where Shohet sees three main risks: operational downtime *(which was seen in the cyber attack reported at a still-unnamed German steel mill in 2014)*, process manipulation *(as seen, for example, in quality control issues at candy-maker Mars, and also in the problematic production of infant formula without necessary nutrients)*, and, finally, compromise of sensitive information *(which the Dragonfly campaign accomplished, in addition to its establishment of a beachhead for potential sabotage)*. Production environments are also in principle susceptible to ransomware.

SCADA Fence is a passive network monitoring system. It provides visibility into industrial environments by performing deep-packet inspection of commands crossing industrial control networks. Like other new-breed detection systems, it uses machine-learning algorithms to establish a baseline profile of normal behavior, and then looks for the anomalies that stand out against that background.

The alerts it generates are intended to give plant managers the information they need to cue a response, and Shohet believes they've successfully reduced the false alarm rate.

Fortscale—Second-generation user behavior analytics.

Fortscale, with headquarters in San Mateo and its research and development center in Tel Aviv, straddles the US and Israel. We spoke with CMO Kurt Stammberger about the company's approach to second-generation user behavior analytics. Their focus, he said, is on helping analysts find insider threats.

Stammberger sees two basic approaches to running enterprise security: permissive or restrictive. A restrictive approach chokes productivity, and it also chokes innovation, because real gains in productivity tend to come when people interact with systems in unexpected ways. But obviously permissive rules carry their own risks, and operating permissively requires that an enterprise have ways of detecting anomalies quickly. Fortscale provides that capability.

We're seeing a market dynamic at work, Stammberger claimed, in which companies try to replace the SIEM. "That's vendor overreach." Fortscale sees the SIEM and an anomaly detection systems as, effectively, two distinct houses of government. The SIEM is the executive. Fortscale's UEBA solution focuses on access and authentication logs, analyzing them to produce the lowest rate of false positives. The solution is application agnostic, and can accept access and authentication logs from an indefinitely large software universe.

We asked him how bad actors might evade detection through behavioral analysis if they were aware that such detection was running. He thought this would be difficult—"an insider threat by definition is abusing an access and authentication rule." The solution looks for anomalies in the use of credentials, then for second and third derivatives of such anomalies.

Stammberger pointed out that what counts as actionable intelligence, or a critical alert, is going to be highly context dependent, and will be different for every organization. Fortscale's system

teaches itself to account for this: behavior that could seem very dangerous in one context may be entirely benign where it actually occurs.

Fortscale's solution is a machine-learning system, and, Stammberger claimed, it's very selective. "It's initially very quiet as it learns, and it's designed to reduce false positives. It establishes user profiles from access and authentication logs to create its own baselines and its own behavioral peering." He emphasized that peering is based on observed behavior, and not on the imposition of an artificial organizational construct. Thus it determines peers by what users do, not by their position in an organization chart.

Stammberger said that Fortscale was growing fast, with some initial major customers and a focus on the Fortune 500 as its market. In principle, their solution can be integrated with a customer's existing systems. "The first question we ask is, 'do you have a SIEM?' If not, then that prospective customer isn't ready for us."

the cyberwire

editor@thecyberwire.com
www.thecyberwire.com

@thecyberwire
+TheCyberWire

About The CyberWire

The CyberWire is supported by CyberPoint International and its community partners. We invite the support of other organizations with a shared commitment to keeping this informative service free and available to organizations and individuals across the globe.

Revised 3.23.16