2017 TAG CYBER SECURITY ANNUAL VOLUME 2

Cyber Security Industry Luminary Interviews

Expert Advisory Research

Dr. Edward G. Amoroso
Chief Executive Officer, The Amoroso Group (TAG Cyber)

Version 1.0 – September 2016

Designer – Vision Creative
Finance – M&T Bank
Promotion – Braithwaite Communications
Administration – navitend
Research – TAG Cyber LLC
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 2017 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 1, 2016

To the Reader:

The interviews contained in this volume were written in the summer of 2016, mostly based on discussions between the interviewee and myself – often in person. They are designed to recreate the stimulating discussions I enjoyed during the research associated with this 2017 TAG Cyber Security Annual. I found these interactions to be amazingly useful and insightful on cyber security, and I hope you feel similarly.

My purpose in including these interviews was to allow the voices of these fine cyber security technologists – *luminaries*, really – to be experienced in a more intimate setting than just hearing them at a conference, or in some magazine or newspaper interview designed for non-technologists. The biggest challenge I had was editing down the material to something manageable. Several of the original interview transcripts ran over twenty pages, and I knew this would be too much. Roughly fifty or so interviews at about three pages-per-interview, resulted essentially in a decent sized book – and I was honestly surprised at the resulting volume. I did the best I could to keep each discussion succinct, and to try to capture the essence of each interviewee's voice and main message.

While the interviewees ranged in their backgrounds – from technical, to marketing, to business, and to even hacker – I was struck by how similarly they all viewed their primary job. Every single one of these individuals and their teams emphasized that their common mission is to support the Chief Information Security Officer and enterprise security team. Because it is the CISO teams in industry and government, every luminary agreed, who are the true heroes of the cyber security wars – not the hackers. These unsung professionals toil day-in and day-out trying to keep our power working, our communications connected, our planes and trains running, and our food and water fully stocked. This is why, I quickly learned, the people interviewed in the volume do what they do for a living – and I found it inspiring.

Pick and choose the interviews as you see fit – they are organized into alphabetical ordering for nothing other than tidiness, an approach that I learned as someone named "Amoroso," would be perfectly acceptable for the alphabetically advantaged few (Nice job, Agari and AlienVault). The interviews are intended as a resource, and each discussion is context-free of others, so you can read them in any order. In fact, I removed the pages numbers, because that would just have reinforced the incorrect view that there is some meaningful ordering here.

I hope you like these interviews.

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

Establishing Trust in Your Inbox

Reducing the Security Risk of Phishing Attacks in the Modern Enterprise

Patrick Peterson, Founder and Executive Chairman of Agari

Every cyber security expert agrees that the most insidious attacks today start with some sort of email-based probe. Links to infected sites, social engineering to extract money, or payloads carrying malware slip through gauntlets to the target user's PC. When the user clicks innocently, this leads to a series of steps including infection, lateral movement, and data exfiltration. As such, the ability to advance payloads into the enterprise in order to defraud partners or customers through fraudulent means continues to nag organizations doing business on the Internet. Luckily, effective cyber security techniques do exist that are standards-based, and that can reduce the security risk of fraudulent email, hijacked domains, and other techniques such as spear phishing popular among the offensive community.

EA: Many CISO teams consider phishing attacks to be essentially unstoppable. Are they wrong?

PP: With so many companies getting successfully phished these days, it sure might seem that way. In fact, it's hard to remember a recent advanced persistent threat attack that did not start with maliciously spoofed email, combined with some sort of social engineering or phishing. The good news, however, is that the security risk associated with email infrastructure can be reduced significantly. One of the most potent building blocks in this regard is an open standard introduced in 2012 called DMARC, which stands for Domain Message Authentication Reporting and Conformance. DMARC enables global visibility of domain name use in email, allowing email senders to better authenticate their identity to email receivers, and to therefore prevent bad guys from spoofing domains in malicious email.

EA: Is DMARC different from previous standards like DKIM and SPF? And what do customers need to do to support such standards for email security?

PP: Good question. Actually, DMARC extends the existing standards you refer to, namely Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM). The result is an improved approach that helps senders establish some important properties. First, DMARC connects the actual identity of an email sender with their sending domain. Second, DMARC provides for the publication of policy-based options, such as message rejection, for how receivers should handle email that is not considered legitimate. And third, DMARC supports real-time intelligence gathering across the Internet for domains of interest, which is the basis for Agari's solution offerings. Some say that DMARC is to email, as Secure Sockets Layer (SSL) is to Web eCommerce.

EA: How does a CISO team use DMARC to improve their email security posture? Are there many steps involved and is there a learning curve?

PP: It's not that complicated. First, the domain owner should publish a DMARC record. For larger organizations, there can be hundreds or even thousands of legitimate servers sending email on behalf of your organization. The challenge is identifying and publishing a DMARC record that accurately reflects all of those authorized email senders so that you can set the policy to reject and block the untrusted senders of email. With many companies, the number of third party services and domains sending email is in a constant state of change and the DMARC policy needs to be monitored and maintained. If the local staff is not sure how to do this, companies such as Agari can help. Second, email authentication should be deployed via DKIM and SPF. This requires publication of records that describe the servers authorized to send on behalf of an email domain. This might also require email servers to be configured to insert DKIM signatures, and yes – companies such as Agari can provide assistance for customers who might need reassurance that they are doing this properly. Finally, the team must enable Identifier Alignment, which is really how the DMARC-supplied aggregate feedback allows identification of where domain identifiers do not align with the email domain.

EA: How do companies like Agari participate in this process? Do you have a specific product or service that enterprise customers would buy?

PP: Yes. Our team offers a cloud-based SaaS solution called the Agari Email Trust Platform, which protects three billion of the world's inboxes from threats such as phishing, targeted attacks, and business email compromise. Agari secures the entire email channel for customers, employees, and partners from advanced email threats. The Agari Trust Network creates a model of trusted email by analyzing an organization's inbound email and outbound email senders, and correlating this information with analysis of billions of email messages per day from the world's largest email providers including Google, Microsoft and Yahoo. Then, that trust

model is used to categorize and prevent untrusted email from reaching the inbox of employees, partners, or customers.

EA: Establishing sender authenticity seems like such an obvious and important technique. Why do you think the industry has not been more aggressive in making this an absolute priority, especially in government applications?

PP: I think there are a couple of reasons. First, DMARC is relatively new, having been established only in 2012. And while it takes time for email receivers like Google, Yahoo, and Microsoft to fully adopt the standard and to begin authenticating emails based on DMARC, we have reached critical mass with more than 90% of inboxes in the United States authenticating based on the standard. This places us in a subsequent phase, where email senders such as enterprises and government organizations need to publish DMARC records and move toward a reject policy. With more and more email being sent through third-party services, organizations also need help by identifying, managing, and maintaining governance over their email senders to enable publishing a DMARC reject policy, which is one of the ways Agari helps customers. A second issue has been with compliance auditors, who are typically less familiar with new standards such as DMARC. Hopefully, the regulatory and compliance community will familiarize themselves with these sorts of controls so that DMARC can become more uniformly applied as part of compliance initiatives. The approach, as you would think, works best when more email senders and receivers take the time to get properly set up. As with the 2005 FFIEC two-factor authentications compliance requirements for Internet Banking, compliance auditors need to set standards that require publishing of DMARC records to protect the public and email communication in general.

EA: We've seen a dramatic increase in targeted attacks such as spear phishing and business email compromise from FBI statistics. Why don't existing protections stop this, and what can organizations do to protect themselves?

PP: Email continues to be the primary way cyber criminals infiltrate enterprises. As much as 95% of cyber attacks and data breaches use spear phishing as the initial entry point. Existing solutions such as secure email gateways and advanced threat protection do not fully protect organizations from targeted email attacks, spear phishing, and business email compromise, because they focus on detecting malicious content and bad behavior. Attackers can evade detection by crafting socially engineered email attacks with no malicious code or URLs, and they can impersonate trusted senders such as internal employees, partners or vendors. The only way to fully stop these types of attacks is to identify the trusted sender of the email, and to focus security controls on defining trusted email behavior in order to prevent untrusted email from reaching employee inboxes. The Agari solution focuses on exactly this type of trust modeling as an additional security layer to protect enterprises from advanced email threats.

EA: *As the founder of a successful cyber security firm, you have an interesting vantage point into present and future trends. What predictions do you have for Internet security in the coming years, especially as they relate to email fraud?*

PP: Actually, I see things getting better, primarily because enterprise security teams and experts are getting better. Sure, the offensive community has surged ahead of the defense recently, as evidenced by one attack after another, including the terrible break-in and theft at the Office of Personnel Management (OPM). But with international government organizations and small, medium, and large businesses beginning to work together using common standards – and DMARC is just one example, I think it is possible for cyber security protections to catch up with the attackers and this will result in a safer and more secure environment for business, government, and industry.

Unified Security Management for Enterprise Cyber Security

Combining the Best of SIEM, Log Management, and Cyber Threat Intelligence into a Unified Defense

Roger Thornton, CTO of AlienVault

The earliest log management tools were designed to reduce the manual burden of having to sift through reams of printed or archived audit trails in support of various compliance and other business objectives. Such work might have been done with mainframes, and it was rare that anything of real security consequence was actually detected. This crude early task has now evolved pretty dramatically into the use of modern security information and event management (SIEM) platforms supporting real-time enterprise data collection, cyber threat intelligence gathering, security event monitoring, and advanced assistance for the threat hunting task being performed today in security operations centers around the world.

EA: Do most companies now have a SIEM – or do you still see gaps in coverage, perhaps in smaller organizations?

RT: Just about every company in the world has some requirement for security visibility and monitoring today. In spite of that broad mandate, I see lots of gaps in coverage for SIEM – and that's not just in smaller companies. It's in most companies where teams are stretched thin and budgets are tight. In most of these companies, you will find a smattering of point solutions. They may operate intrusion detection at the perimeter or some vulnerability scanner collecting basic logs, but very few have all the controls in place, or have them integrated for effective security monitoring and analytics. A traditional SIEM is a data aggregation platform, and that aggregation can be a difficult chore in itself. But it is a pointless one if you don't possess the pertinent data to perform the required enterprise cyber security analytics.

EA: What are some of the technical challenges involved in trying to unify the various tasks related to log, event, and incident analysis?

RT: Long gone are the days that any single point solution can effectively detect threats. Detection now requires continuous monitoring of multiple facets of an infrastructure, which will require multiple different controls to be deployed and integrated into a platform that will support analytics. The first challenge is collecting data – and to do this, you must know what's on your network. You also must know the vulnerabilities across your network assets, and this information must be up to date, rather than a snapshot in time. You will also need detailed information about the behavior or your network and the systems running on it, including what protocols are present, what connections are being made, and what users are doing. Finally, you need detailed information about the threats themselves – and this changes constantly. Once you have all that, you then need to pull it all together into an analytics platform so that you can find the bad guys with enough accuracy to direct action, but without too many false alarms. Our approach at AlienVault is to address this complexity through an integrated solution that orchestrates data gathering, security controls, threat intelligence, and analytics into one simple, easy to use package.

EA: What role does accurate threat intelligence play in deriving good intelligence from collected data?

RT: Well, I can tell you that *bad* threat intelligence involves stale indicators of compromise (IOC) in the form of virus signatures, URLs, domain names, and IP addresses. Such stale data is often useless, because the attackers have moved on, and the IOC may now be pointing at the wrong source. This threat intelligence can be improved by adding context, and to do so, requires that it be constantly updated. This work can be intensive, so you will either need a good threat feed from a vendor or sharing community, or you will need an in-house security research team. The *best* threat intelligence can be consumed directly by security controls to produce effective preventive or mitigating action. Instead of just providing IOC's, the best threat intelligence provides the specific tuning rules for your IDS, vulnerability scanners, firewalls, network analysis tools, and SIEM correlation engine – and this is the approach we try to take at AlienVault.

EA: Do you have opinions about open source tools in enterprise? What factors should a CISO team take into account before downloading and using an open source tool?

RT: You would be hard pressed to find a company that does not use some open source tools within their information security program. And we all know that attackers make use of open source tools in their exploits. With that said, there is certainly a place for both open source projects and commercial products within any security team – albeit with the time and expertise required to make open source truly work. At AlienVault, we maintain an open source project called Open Source

SIM (OSSIM). This open source offering, like our commercial Unified Security Management product line, is an integration and orchestration platform for a collection of embedded security tools. OSSIM integrates several open source security tools including snort, nmap, and OpenVas. Like most open source tools and projects, OSSIM works best when used by experts and researchers with deep security skills.

EA: *How do you see virtualization of the data center and evolution of the network to SDN as affecting the unified security management task in the modern enterprise?*

RT: Because of the rapid adoption of virtualization, products that are monolithic, expensive, hardware-centric, and bound to single operating environments are simply doomed to extinction. Virtualization allows for segmented computing into fine-grained processing regions, which significantly reduces the attack surface at any one point. Furthermore, virtualization and SDN can unify security controls, provided they are designed to support such action. Virtualization and cloud environments provide templates that will greatly simplify the setup and configuration of security tools. And the promise of SDN for security is the ability to put monitoring and analytics agents at just about any point in the network with the same ease required to deploy software into operating environments. At AlienVault we accomplish this through software sensors that can be provisioned with minimal effort into virtual data centers and cloud environments.

EA: *Any trends you're seeing in the threat and attack space? Is it getting harder to detect attacks due to increased offensive capability – or is it getting easier to detect attacks due to more automated, feature-rich tools*

RT: Attackers have an enormous advantage, because they can decide when, where, and how to strike. But once they've made initial inroads, then they are like strangers in a foreign place. They have to search the environment, move laterally, attempt access, and try to exfiltrate. The truth is that identifying this behavior is actually quite easy, as long as you have the right data, tools, and knowledge of what to look for. So, while it may be impossible to keep someone out, the good news is that it can be relatively easy to catch them, once they are in. Security teams should thus make sure to invest sufficient time and money into getting good at threat detection and incident response. Large companies with sophisticated security teams have been doing this for years, and now it's catching on with everyone.

IT Infrastructure Utility as Basis for Enterprise Network Security

Offering Utility Services From Layer 4 Provides an Underlying Foundation and Cyber Protection Model

Bruce Flitcroft, CEO of Alliant Technologies

The traditional IT model has been challenged in the past for its lack of sufficient flexibility to deal with the high level of change inherent in enterprise applications such as workflow, office applications, and databases. By offering IT Infrastructure as a utility service at layers 4 and below, the advantages of metered, on-demand usage can be combined with flexible interfaces to application level IT. The implications on enterprise infrastructure security are also meaningful as network level protections can be embedded into a utility model, with requisite alerts and alarms exported from a defined utility interface.

EA: Let's start with a broad question. What is an IT Infrastructure Utility?

BF: An IT Infrastructure Utility delivers information technology and networking capabilities to an enterprise in a way that is similar to how power companies provide electricity. Modern enterprise IT, network, and security teams tend to waste a high percentage of their budget just trying to keep the lights on in their infrastructure. This detracts from their ability to focus on delivering flexible, high-value capabilities to their users. At Alliant Technologies, we've partnered with global service providers like AT&T Partner Exchange and world-class technology providers such as Cisco to create an IT Infrastructure Utility that has many advantages. For information technology and network teams, it simplifies management and reduces cost. For enterprise security teams, it ensures that infrastructure security processes like patching, vulnerability management, and security operations center (SOC)-based oversight are done in a timely and accurate way.

EA: Does an IT Infrastructure Utility operate across the entire range of network and IT services?

BF: It certainly supports all layers, but we've found that the best approach is to focus the IT Infrastructure Utility on network layers 4 and below – as defined in the OSI stack. This allows the enterprise to focus on layers 5 and up, which includes the application layers that are seeing so much more SaaS-based deployment these days. For security teams, this approach creates a flexible, virtualized perimeter of the enterprise's private domain, thereby extending the enterprise's private infrastructure to absorb off-premises locations, including cloud services, and also providing for comprehensive security management and compliance. Furthermore, the IT Infrastructure Utility collects all the foundational operating data, including syslog and NetFlow, and makes this data available to higher level enterprise tools such as the SIEM, in the form of log file access or alarm tickets.

EA: Can you say more about the impact of an IT Infrastructure Utility model on security? Are the interface seams between utility services and the upper level applications possible points of attack?

*BF***:** An IT Infrastructure Utility model provides a firm foundation on which security can be built and managed. This is accomplished through the use of standard reference architectures, which reduce the number of vendors and configuration variables that need to be pre-validated, thus *decreasing* the number of possible attack vectors. In addition, the network infrastructure itself is becoming a security sensor, and it needs to be maintained to perform that function. Security is integrated into the reference architecture in a way that allows enterprises to select how they want to enact and enforce security policy without dictating a particular method. Furthermore, the IT Infrastructure Utility is proactively managed, which ensures that the infrastructure stays up to date with vendor security alerts.

EA: Would the IT Infrastructure Utility service provider become part of the enterprise incident response team for issues that occur at layer 4 and below, such as DDOS attacks?

BF: Yes, the IT Infrastructure Utility plays a vital role in incident response, especially where security remediation and recovery require action at layers 4 and below. During an incident, the utility provider makes reactive changes to device configurations based upon enterprise-provided mitigation strategies specifically related to the deployed security configurations. And the proactive configuration management in an IT Infrastructure Utility also reduces the attack surface by eliminating security issues as they are discovered by manufacturers and ensuring that software and configurations up to date. Additionally, IT Infrastructure Utility provides valuable operational and performance data to the enterprise SOC in support of threat analysis.

EA: How does an IT Infrastructure Utility accommodate the one-off needs of an enterprise buyer? This is a big issue for security teams.

BF: Our primary approach taken at Alliant is to accommodate the one-off needs of the enterprise through flexibility in the architecture and support for as-needed changes via proactive management. This is not the same as letting the buyer dictate the equipment delivered in the utility service, but it does allow for a complete set of features – and this includes enterprise security features – to be enabled in the infrastructure. A good example is that an IT Infrastructure Utility can accommodate real-time changes in the face of an active attack. This can be viewed as one-off support, but we see it as a normal component of a live response.

EA: Any final thoughts on the IT Infrastructure Utility services for enterprise security teams?

BF: Our mission is to bring the utility model to private, on-premises network infrastructures, and this has good implication for security teams. Take risk management, for example – an issue that is well understood in the security community. Historically, there has been an absence of risk sharing in the traditional IT Industry. The enterprise bore all the financial, technical, deployment, and operational risk. The enterprise had to piece together the different technology vendors required for an end-to-end solution, and was responsible for making sure the various parts worked together. In this traditional IT model, the enterprise was responsible for coordinating deployments across vendors while assets lay idle. This risk model changes with the IT Infrastructure Utility, which involves sharing of the enterprise risks, and this includes security capabilities at the lower layers. As the industry gains awareness that this IT Infrastructure Utility is available in the market along with its financial, operational, and security benefits, enterprises security teams will come to recognize the value of this approach for protecting the enterprise from cyber attack.

Recruiting CISOs and Building Their Teams

Using Retained Search as a Risk Reduction Technique for the Modern Enterprise CISO Management Team

Joyce Brocaglia, CEO of Alta Associates

Few technology disciplines have seen as much change in career management than cyber security. Born as a sleepy technical focus in the 1980's, the closest an early security professional would come to the C-Suite would be if some senior manager lost a badge. Since then, however, the profession has begun transforming into an essential component of every organization. In some cases, the Chief Information Security Officer now serves as a valued member of the executive team. But with this new role for enterprise security professionals comes the requirement to learn new skills in recruiting team members, training new staff, and managing one's care. Alta Associates has been recruiting information professionals since 1994, when they built the first ever information security organization as a result of a Russian hack.

EA: Joyce, you've been at this for nearly three decades as Alta's CEO. You are certainly in the best position to explain how retained search works for an enterprise.

JB: Retained search is a partnership between the executive search firm, enterprise security team, and their HR business partner. At Alta Associates, we are brought in to partner with an organization that is either looking to hire their first time CISO, elevate the status of their current cybersecurity or risk leader, or build a world class team. With our deep subject matter knowledge, we begin searches by educating the organization on the current competitive landscape, including the market value of candidate's compensation. We can also provide feedback and guidance on where a candidate or role should be placed on the organizational chart, as well as how to develop the job description and create a successful interview process. Because we have been specializing in recruiting information security teams for so many years, our deep trusted network allows us to be an effective partner. Also, our recruiters

have spent years of their careers specifically focused on cybersecurity. They are all subject matter experts who understand the nuances of the roles, the particular industry specialties, and the culture of various organizations. We believe that such knowledge, coupled with their deep trusted relationships, make Alta second to none when it comes to delivering top talent. We begin our partnership with clients by initiating a launch call with the hiring manager, key stakeholders, and HR partners. From that call, we finalize the job description, create an outreach strategy for candidates, and set a timeline of weekly follow-up calls where we present resumes, gain feedback, or create offers. Throughout this process, the client works with a relationship manager and team of recruiters who have done an extensive search and in-depth candidate interviews. By the time an offer is made, the client is confident that they have done their due diligence in interviewing the best and brightest for the role.

EA: Do you see any trends in cyber security search? For example, is it getting easier or harder to find competent, trained professionals for clients?

JB: It's getting harder, and all the research supports that conclusion. In the past five years, the demand for Infosec professionals has grown over three and a half times faster than other tech roles. Job postings were up 74% and an (ISC)2 workforce study showed that 62% of respondents said their organization had too few Infosec professionals. This is true up and down the line, in terms of experience and expertise, but is even more intense for senior positions. The role of the CISO is becoming more and more complex. Not only do CISO's have to be technically competent, they have to understand the regulatory, privacy and risk implications and impacts to their organizations. As if that's not enough, they have to have the business acumen and communication skills to convey technical solutions to the board, audit committees, and key stakeholders. It takes experienced and knowledgeable recruiters to discern those qualities.

EA: What are the advantages of bringing executives in from the outside? Does this have drawbacks?

JB: The advantages are numerous, including the fresh perspective, optimism, and diversity of skills that executives bring from positions held outside the hiring organization. It cannot be underscored how critical such diversity of background can be to create innovative programs. Research shows that eclectic groups of people with different backgrounds, gender, ethnicity and training are more productive and innovative. The drawbacks associated with recruiting executives from the outside relate to the investment of time and energy that must be allocated to the process. That's why it's important for hiring managers to take control of the process, not leave it up to the HR team, which is typically already overburdened and lacks an understanding of the nuances of the roles and the competitiveness of the market. It's imperative for hiring managers to make recruiting a priority on their calendar and partner with a specialty search firm that can drive and manage the hiring

process. Staffing and recruiting are not outsourced functions that allow hiring managers to remove themselves from the process.

EA: I know you are passionate about empowering women in cyber security. What are you doing to improve the posture and presence of women at all levels?

JB: As you know, in 2002 I founded the Executive Women's Forum (EWF) on Information Security, Risk Management & Privacy. Today we are the largest member organization dedicated to engaging, developing, and advancing women leaders in our field. The EWF hosts events and programs throughout the United States, including regional meetings, networking diners, and informal gatherings. EWF membership includes women at all phases of their careers, from staff positions through the C-suite, and our programs are developed to uniquely help them at each stage. We are probably best known for our annual conference that gathers over 400 women thought leaders in our industry together. Our 14th Annual National Conference is entitled "Balancing Risk & Opportunity: Transforming Cybersecurity, Risk & Privacy Beyond the Enterprise." To give you an idea of our programs, we have a mentoring initiative for staff women. We provide a Leadership Journey, which is a comprehensive leadership development program for middle managers. And we sponsor a Women of Influence Round Table for our most senior ranking members. Over the past decade and a half, I have also seen great strides taken by *men* in the field who are recognizing the importance of diversity of thought and are taking positive steps forward in their efforts to hire and develop women on their teams. These men have to act as role models and encourage their peers to do the same.

EA: How do cyber security professionals find a search firm if they are interested in making a personal career change?

JB: They should work with an executive recruiting team they feel comfortable with. Years of experience and past performance should be taken into account when establishing a relationship, but there must also be a personal, human connection. People do business with people they trust. Time and again our clients and candidates tell us how much they appreciate the time that we take in understanding their needs and goals and how diligently we work on their behalf to meet their personal objectives. We really enjoy our work and the people we work with, and we take our responsibility in bettering people's careers, teams, and the industry as a whole very seriously.

Measuring Security In Mobile Apps

Helping the Enterprise Quantify and Understand the Risks of Downloaded Apps

Paul Stich, CEO of Appthority

One of the first and most memorable warnings computer security experts offered in the humble early days of the industry was that PC users must be careful about downloading software from unknown media such as floppy disks. Given such early emphasis, it is hard to explain why so many current mobile users, even ones employed in critical infrastructure settings, think nothing of downloading software from an app store with little understanding of whether Trojans and other security vulnerabilities might be present. It seems obvious that security measures and risk reduction measures will be needed in the coming years to rectify this trust issue for mobile apps.

EA: Should people trust the apps they download to their mobile devices from popular app stores?

PS: No, they should not blindly trust downloaded apps. And this is tough, because we all know that apps are the new form of software consumption. In the past, users would purchase or subscribe to software at the point of download or sale. The developer would get financial compensation directly from the end-user who purchased the software. In the new app-centric model, however, most apps are downloaded for free, yet developers are still expected to produce a great product, provide updates and additional content, and provide support – all for free. Users have to ask themselves where the catch in all this might be. Because developers are not being compensated for their work, they are incentivized to harvest user data and then sell that data to data brokers, advertising networks, and other third parties. Thus, for the most part, when an app is free, your personal data is the product.

EA: Do you see differences in the way Apple and Google review, and then approve or disapprove, mobile apps in their respective stores?

PS: In the past, there was a more stark difference between the app review processes at Apple and Google. Apple relied mostly on human review with very long review cycles and paid close attention to adherence of the App Store's terms, conditions, copyright issues, and app functionality. Google, in contrast, relied mostly on very brief automated app analysis, and focused on identifying malware and other malicious app activity. Apple has since shortened their app review process, likely through automation, and Google has added human review of copyright concerns and potential violations of Google Play's terms and conditions. Nevertheless, Apple still takes longer to review and approve apps, and still has a much higher bar in terms of acceptance criteria. For the most part, both do an adequate job of preventing malware from entering the store. Apple does a better job of preventing cloned or fake apps from entering their store, but neither does a good job of preventing apps that, while okay for personal use, are often deemed too risky for enterprise use. This is because so many apps are riddled with code level vulnerabilities, thus demonstrating data leaking behaviors, which are designed to harvest and share user data.

EA: Are companies experiencing significant mobile app breaches?

PS: For the most part, companies are not even monitoring their mobile devices, so they are not detecting mobile breaches. MobileIron, for example, disclosed that less than 5% of their customers have an anti-malware solution installed on their mobile devices. It's safe to say that all companies are experiencing minor mobile breaches, like apps stealing address books, calendars, and other device information which contains sensitive corporate data. However, a major mobile app breach has not yet been disclosed. For now, it can be seen as "death by a thousand data leaks". Keep in mind that each small data leak, even when it is a leak of personal data, can, in the future, enable a much larger non-mobile hack in the way of a targeted phishing attack, for example.

EA: What techniques have companies used in the past to evaluate software security? Do these still apply in the mobile context?

PS: In the past, companies built and managed whitelists and blacklists to determine what software was acceptable for use within corporate environments. Although some companies have tried to replicate this model in mobile, they've quickly realized that it simply does not scale in the mobile context. There are millions of apps out there, and they are versioning so quickly – sometimes more than 10 times per year – that building and managing a whitelist manually is impossible. Automation is key to not only quickly analyzing the apps (as opposed to long pen-testing approaches of the past), but also to automatically remediating against non-

compliant apps and devices. Rather than blacklisting apps, for example, many enterprises are now focusing on blacklisting certain risky app behaviors.

EA: How hard is it to review an application to quantify the risk? Do you observe its behavior or do you review the binaries?

PS: Traditional software security analysis usually required access to source code. But source code is not available for public apps on employee devices. Traditional analysis also often involved long drawn out penetration testing processes, where a researcher would try to find weaknesses in software by running different commands and testing the software through different scenarios. In the mobile world, however, customers need to identify app risk as quickly as possible, given that employees are downloading new apps, or new versions, every day. Because of the high number of third party software like SDKs and libraries being used in apps, relying heavily on static analysis of the app binaries can also lead to a lot of false positives. Thus, it is essential to leverage dynamic analysis, as we do at Appthority, to see how apps really behave at run time. Using an instrumented sandbox, analysis engines are able to track app behaviors and create an adequate risk profile of each app version.

EA: When you perform these reviews, what models are you using to establish risk? Do you fold real-time risk intelligence into the analysis?

PS: Our research team monitors the threat landscape for new types of attack vectors or new evasion techniques by malware families. The team then writes rules for our analysis engines to execute during run time app analysis. From there, our engines look for these risky behaviors at scale across millions of apps to identify the prevalence of these app behaviors, as well as identify patterns that could indicate combinations of behaviors that form a behavioral signature for emerging threats. Because the analysis engines monitor behaviors at run time, there is real-time intelligence on not only how the app behaves, but also where the app is communicating. A real time URL/IP threat intelligence feed is then overlaid with the analysis results to see if additional threats are identified.

EA: Have there been some spectacular finds that you've seen in your years now at Appthority? Maybe you found some app that had a ridiculously dangerous Trojan?

PS: We've seen a lot of dangerous Trojans in the wild, mostly on third party app stores. In these cases, malicious actors inject malware into an otherwise good app and release it to unsuspecting users. However, one of the most notable Trojans was a phishing attack in the official Apple App Store. In this case, the developer of a legitimate app accidentally used a contaminated third party SDK that included malicious code to prompt users to enter their Apple ID, only to intercept it and later use it to attack users, often with ransomware.

EA: Paul, you're an industry veteran in cybersecurity. What changes have you seen in the past few years on both the offensive and defensive sides?

PS: There are over 1.5 million cyber security job openings just in the United States alone. With the huge shortage of qualified IT and security professionals within the workplace, corporations continue to struggle, not only with getting increased budgets to purchase security solutions, but also with finding competent staff to manage and use all of these tools and services. As a result, there has been a big push for automation of tools and services that can run on their own and make the IT and security teams more productive.

Assuring Communications Through Smart Scrubbing

Using Visibility, Intelligence, and Algorithms to Protect Networks From DDOS and Advanced Threats

Eric Jackson, VP Arbor Networks

Distributed denial of service (DDOS) attacks have always been particularly insidious, simply because with only modest resources on the part of the attacker, great havoc can be brought on targeted victim assets. Traditional botnet-based DDOS attacks have been successful by creating enormous volumes of information that are tossed at a gateway. The problem is that as ISPs and security vendors have gotten much better at filtering these volumes, the attacks are getting smarter and arguably more dangerous.

EA: Eric, in terms of layer 3 packet volume, what are the largest DDOS attacks you've seen?

EJ: We've seen denial of service attacks with packet volume sizes up to 550Gbs – or at least that's the highest size attack I'm allowed to mention. To get to this size, the attackers launch several varieties of UDP reflection and amplification DDOS attacks, including DNS, SSDP, and NTP. With more sophisticated attackers, we often see multiple DDOS vectors used simultaneously, and the smart attackers monitor the efficacy of their attacks, and will adjust based on the defense. Also, with reflection and amplification, the attacker spoofs the source IP of the target, pretending that he's the targeted system. The attacker will then send lots of query packets to hundreds, thousands, tens of thousands, or millions of misconfigured servers exposed on the Internet. Once those misconfigured servers receive the spoofed query packets, they turn around and dutifully answer the intended target of the attack, pummeling the target IP addresses, and often filling up peering links and downstream transit links due to sheer attack volume.

EA: Do you think that better and more uniform deployment of source address verification is a solution to the DDOS problem?

EJ: Yes, more comprehensive source address verification will certainly help, but in the meantime, the typical enterprise security teams will need a comprehensive protection strategy. And no organization, including the largest ISPs, can deal with these attacks alone. Every enterprise needs to put a security plan in place, which would seem like such an obvious step – and yet, Arbor produces an annual Worldwide Infrastructure Security Report (WISR), which finds year after year that most survey respondents have no plan. Another important aspect of successful DDOS defense is having visibility and insight into traffic. Any Internet-connected organization should be able to detect, classify, and trace DDOS attacks to their entry points. And finally, it is essential to have timely DDOS mitigation assistance in place organically or from a service provider.

EA: What about application layer DDOS attacks? What are you seeing now, and what trends are you anticipating?

EJ: These attacks succeed because the victims are vulnerable and unprepared, often running Internet accessible software applications that are fragile and poorly designed. We see, for example, application layer attacks against badly administered Web servers. We also see attacks against SSL/TLS termination points that exhaust the ability of the server endpoint to handle legitimate incoming encrypted sessions. And we see DNS as a popular application layer target, as evidenced during the 2012 attack campaign against US banks. More recently VoIP services are being exploited based on Session Initiation Protocol (SIP), where SIP call-control floods are being launched to tie up SIP PBX and Session Border Controller systems to block handling of legitimate voice calls. In the online gaming sphere, we're seeing more and more attacks against gaming networks and services, which involve reverse engineering of the application layer protocols specific to an online game, followed by attempts to overwhelm the game servers with streams of legitimate looking requests.

EA: What about trends in the motivations behind these attacks? Are DDOS attacks today mainly just kids playing, or is there a more serious motivation?

EJ: In our most recent WISR, ideologically motivated DDOS attacks topped the chart for the first time ever. This is bad news, because defenders know that ideologues are very persistent. Attacks sparked by online game disputes are also common, along with extortion, botnet testing, and vandalism. And many DDOS attacks are now launched to *distract* security teams from observing data exfiltration or online fraud that might be going on during the DDOS attack. This is an important reason for security teams to have plans in place, because any distraction during an attack can divert resources away from the most critical risk. So, as you can see, the motivations behind DDOS attacks can certainly vary.

EA: Are you seeing crossover between botnets being used for espionage, fraud, and DDOS?

EJ: Increasingly, we see botnet designers adding DDOS capabilities to botnets, which had previously been used for other things. We've seen, for example, an especially disturbing crossover between ransomware and the DDOS world, with ransomware coders starting to add DDOS capabilities to their malware. This is a troubling new trend because it creates the possibility of an internally generated DDOS attack springing up from malware getting into the enterprise.

EA: That sounds like a huge, huge problem. Are most enterprises prepared to deal with internally generated DDOS attacks?

EJ: Well, we all saw what happened with SQL Slammer, Nachi, and Blaster in 2003. Organizations suffered serious outages related to these network worms, so this is likely to occur again – if only because most enterprise networks haven't changed much since that time. Fortunately, the same tools and techniques used to deal with outbound DDOS attacks can also apply to internal-only DDOS attacks. And this includes having a response plan in place.

EA: It seems like most organizations tend to ignore the DDOS threat until they're actually under attack. Are there any commonalities between more traditional security focus on disclosure, integrity and break-ins, and the capabilities required to successfully mitigate DDOS attacks?

EJ: Absolutely, and at Arbor, we have evangelized a six-phase model for dealing with any type of security event, whether it's a systems compromise, a data breach, or a DDOS attack: First, there is *preparation*, which involves getting the tools, techniques, processes, procedures, personnel, and training in place in order to successfully handle a security incident. Next, there is *detection*, which involves having the ability to know that an attack is taking place. This is followed by *classification*, which involves determining the attack type, seriousness, and potential implications. Next, there is traceback, which means being able to trace an attack to the source where it is entering the network. This is followed by *mitigation*, which is the ability to stop an attack by filtering out DDOS traffic, quarantining compromised hosts, and disabling breached applications to stop data exfiltration. The final step involves *post mortem*, which is the step everyone forgets. It involves determining what went well during an incident, identifying gaps, and feeding the output of this step back into the entire cyber security process.

Using Deception to Detect Cyber Attacks

Creating Clever Means for Luring Malicious Bots to Reveal Themselves in Real-Time

Tushar Kothari, CEO of Attivo Networks

The use of deception in warfare is well established in dealing with difficult adversaries. The method is particularly useful in asymmetric situations where the defense is at a clear disadvantage. Many military organizations with limited resources, for example, have resorted to using decoy tanks and fake troop activity in an attempt to even an unbalanced playing field. Cyber security, it turns out, is the ultimate asymmetric situation where an enterprise has limited means to stop attacks that can come from nation-state military groups. Deceptive cyber security defensive techniques are thus extremely effective in producing uncertainty in an adversary, while also increasing the chances of real-time detection.

EA: First of all, do you believe that the most advanced attacks can be detected in real-time by an enterprise defense? Or have we entered a new phase where advanced breaches are inevitable and where we must focus instead on response?

TK: It is true that many current enterprise defenses have proven to be unreliable and have been successfully breached by sophisticated attackers. We see the effects of this every day with high profile breaches into companies that clearly have real-time defenses in place. Furthermore, the predictability of current defenses, and the lack of a true security perimeter, help the attackers and reduce the effectiveness of a prevention only defense. This should be no surprise to students of warfare, who know that war is based on deception, and typically won using the element of surprise. Deception turns the table on the attackers and makes the deceivers become deceived. We believe at Attivo Networks that it is the use of deception that will help make advanced attack detection more achievable.

EA: *What have been the challenges of intrusion detection over the years? Is the adversary just that good?*

TK: Part of the challenge is that as the years progress, the network morphs. Just after each organization finishes building a perimeter around their network, for example, business requirements lead to the punching of holes in this perimeter to enable services such as remote VPN and cloud services. The perimeter soon begins to look like Swiss cheese with holes all around. And as if open back doors are not enough, social engineering with phishing and complex mobile device management help leave the front doors open for getting infected, regardless of the strength of the perimeter. With hundreds of reported breaches each year, clearly our networks have become more vulnerable, and a traditional line of defense is destined to fail. Now, you are also correct in suggesting that the adversaries have become better equipped and financed, often by nation state governments. Also, some of the successful attacks have yielded significant financial benefits for the attackers, which has in turn attracted more sophisticated adversaries. And security budgets have not kept pace, so organizations remain challenged to recruit and retain trained security staff to combat these highly sophisticated adversaries. These are tough challenges, but enterprise teams understand the problem, and there is every reason to believe that the situation will improve for the defense.

EA: What is the basis for using deception in cyber security to detect attacks? Is the idea to be stealthy enough that the adversary is tricked? Or are you really just dealing with automated botnets?

TK: The key focus of deception technology is to turn the table on the most sophisticated adversaries. The idea is to deploy authentic and realistic deception, which is indistinguishable from real assets and is deployed in a fashion that makes it irresistible bait to the attacker. The technique is extremely efficient and effective in catching the intruders, who can be deceived and misdirected into a maze of traps and deceptions within the network. Throughout mankind's history, traps have been developed to catch pretty much anything of value. The same concept rings true for efficient cyber security threat detection. The Attivo ThreatMatrix Platform, for example, deceives the attacker into believing that he has succeeded in his attempt, engages with him and after tracking his lateral movement and behaviors, and extracts the valuable forensic information required to stop and derail the attack.

EA: What's been the practical experience for companies using deception to detect cyber attacks?

TK: Our customers have experienced tremendous improvement in their ability to detect the attackers early in the kill chain during reconnaissance and lateral movement. Given resource limitations and staffing shortages, the Attivo ThreatMatrix platform is designed for high efficacy, high fidelity alerts, and automatic ingest of attack forensic information during engagement to accelerate

incident response actions to automatically quarantine infected systems and update prevention systems to block attackers. Deception also employs techniques to deceive and misdirect an attacker since it is signature-less. The solution is well suited to detect zero day attacks from advanced threat actors. The Attivo ThreatMatrix platform provides an integrated sandboxing technology, which has allowed our customers to extract signatures and TTP (tactics, techniques, and procedures) out of most sophisticated attacks including polymorphic malware. Customers have also called Attivo the "eyes and ears of their network," since we can provide early visibility into internal and external threat actors, and we can detect the use of "harder to detect" stolen credential and man-in-the middle attacks.

EA: Do cyber deception methods require the corresponding use of honey pot content to make things realistic?

TK: While honey pot technology is not required for initial detection, it is extremely valuable when you want to engage the adversary for a longer period to extract full TTP. Creating a high interaction honey net is a valuable component of a deception platform – though today's deception technology goes much further than a legacy honeypot approach.

EA: Do you think the day will come when every CISO team uses deception along with other common techniques such as firewall and access control?

TK: An adaptive defense requires a mix of prevention and detection solutions. It is clear that prevention alone doesn't work and that trying to find the needle in the haystack is too resource intensive with monitoring or Big Data approaches. With this in mind, given the efficacy and efficiency of deception, there is no question that deception will play an important role in the security stack. Deception is accurate where other solutions have proven they are not reliable, and it is easy to deploy, operate, and manage during incident response. Deception is also efficient and immune to resource intensive false positives. Plus, the automated attack correlation, forensic reporting, and attack automation inherent in deception, take cost out of continuous threat management and response. At Attivo, we often recommend that our customers actually *have some fun* turning the tables on attackers. It can be rewarding to "deceive the deceivers" by delaying and misdirecting their attacks, which increases their cost of business immensely and gives an organization much needed time to react before damages can be done. Well-designed and implemented deception can lead to the situation where the hunters become the hunted.

Protecting Industrial Control Systems

Building Appliances to Ensure the Integrity of Critical ICS and IoT Devices

Francis Cianfrocca, Chairman of Bayshore Networks

The intense hype around potential cyber security threats to industrial control systems and Internet of Things (IoT) devices and systems may be warranted. The electromechanical and electronic control systems for the critical infrastructure components that make our society operate safely – think power, energy, water, transportation, and communications – are unfortunately vulnerable to a plethora of cyber attacks. Reasons for this include weak legacy design, improper security configuration, weak software training for the engineers, and on and on. Regardless of the reason, few aspects of cyber security seem nearly as important to global society as the effective protection of our industrial control systems from attack in areas of technology known as information technology (IT) and operational technology (OT).

EA: Should we all be worried about terrible cyber attack scenarios to critical infrastructure systems such as power generators?

FC: I know that as an active participant in the industry, I am certainly worried about these types of large-scale cyber attack scenarios, especially to the type of critical infrastructure components you mention, including power, telecommunications, manufacturing, and transportation. Whether these scenarios are something that everyone should be worried about – well, my hope is that we in the security industry can work together to reduce the associated risk so that people don't have to be so worried. Unfortunately, we're all just getting started with this. So much existing legacy equipment, software, and processes are just plain insecure. This leads to a very big cyber security risk in information technology, operational technology, and industrial control systems.

EA: Anyone who attended engineering school knows that industrial engineers are obviously quite intelligent. Why is it that they've gotten the security so wrong?

FC: Cyber security was never an original requirement in the development of most industrial and operational systems. So it should come as no surprise that these requirements were not prioritized. And you're correct that industrial engineers are intelligent. It was never about their competency as engineers, but rather about the respective focus they've had on functional requirements for stopping cyber attacks. Now that everyone knows the industrial control infrastructure, operational Internet, and related systems are vulnerable to cyber attacks, the entire IoT and ICS industry is scrambling for solutions. Our approach at Bayshore Networks has been to focus on the underlying communication protocols that connect devices to the monitoring and control functions usually operated from a management center. The good news is that this approach allows us to fine-tune the policy controls required to prevent malware from causing serious consequences. The bad news is that this is not an easy process.

EA: Is it hard to reverse engineer a legacy ICS protocol? What techniques do the best organizations use to get this done?

FC: At the lowest level, engineers just have to collect the data, analyze the communications, and then help determine what protocol steps are involved. A majority of existing systems luckily use Modbus, so this emphasis simplifies matters, and our engineers have optimized our tools to deal with this protocol. In other cases, we've tried to create generalized solutions based on automation that helps generate policy controls faster and in a less error-prone manner. Keep in mind that this is mostly for legacy systems with protocols designed many years ago. When new systems are designed and put in place, we can do a better job because the security can be designed from the ground up.

EA: What is the best solution to keeping hackers away from control systems? Is it encryption? Firewalls? Monitoring? All of the above?

FC: It would be easy to just answer "all of the above," but that is too simple, and most companies probably cannot afford to do everything they'd like to do. Instead, the best approach is to manage security risk through a comprehensive program of technology, architecture, and process. The Bayshore IT/OT cloud-based gateway was designed to help orchestrate this overall risk reduction. It supports granular content inspection, industrial protocol filtering based on policy, and is applicable to a wide range of technologies including Industrial Automation and Control Systems (IAS), Supervisory Control and Data Acquisition (SCADA) systems, and even smaller programmable logic controllers (PLCs). The goal is to help security teams avoid having to throw a hodge-podge of different security solutions at the industrial control security problem.

EA: Do you see a difference in new ICS applications – in terms of their security – than legacy systems that were put in place many years ago?

FC: Obviously, an ICS designed today is going to have better support for remote administration, native cryptography, access control, proper authentication, and even code security. But the vast majority of existing equipment and software in every industrial setting remains legacy. And this problem is not likely to go away for some time.

EA: What other types of solutions do you envision for protecting industrial systems from attack?

FC: Techniques related to software-defined segmentation are particularly promising, because they remove the weaknesses inherent in a larger perimeter solution. I also like recent advances in identity management, which are directly applicable to IT/OT systems. In particular, contextual, adaptive authentication based on a range of identity indicators is a good direction for protecting ICS devices. Finally, the overall threat intelligence process and ecosystem are so much more accurate and timely, and this includes intelligence about OT protocols.

Enforcing Strong Task Isolation on Endpoints

Preventing Breaches and Identifying Targeted Attacks Using Endpoint Isolation

Simon Crosby, CTO of Bromium

All too often, the familiar endpoint PC continues to represent the softest spot in the enterprise security architecture. One reason for such weakness is the proximity of the PC to that human user – the one that every CISO team member knows is probably making bad decisions about cyber security. Another more powerful reason endpoints have remained so vulnerable is the fragility of endpoint operating systems to even the most rudimentary attacks. Windows computers, in particular, have been notorious for allowing cyber attacks, and this is not likely to change soon. In response to this weakness, new cyber security techniques are emerging that protect endpoints using clever CPU virtualization that can isolate untrusted tasks, thus making the system more resilient to malware.

EA: Simon, as CTO of Bromium, what sort of endpoint vulnerabilities and attacks are you seeing these days? Do viruses, for example, still find their way onto PCs?

SC: Sadly more than 90% of enterprise breaches start with a click. This includes malicious attachments, malware downloads, malvertising, Java usage, Websites, infected media, improper USB use, and infected executables – all punching holes in the perimeter. Conventional "detect to protect" tools fail, both at the network perimeter and the endpoint, because virtually all malware morphs into new, undetectable variants in under a minute, making signatures useless. And the thousand-fold increase in crypto malware signals a shift from manual breaches with stealthy infiltration and data theft, to machine-timescale breaches that can bring an organization to its knees before the first alert. So the answer to your question is that yes, malware still finds its way onto PCs.

EA: So what can the enterprise do to protect their endpoints? Is the protection of PCs a lost cause?

SC: First of all, security teams need to recognize that the traditional "detect to protect" approach using signatures will continue to fail. Moreover, there will always be exploitable application and operating system software vulnerabilities, not to mention foolish users. So these common themes will not go away. Since the typical breach results from the failure to protect a single endpoint, the overall endpoint protection architecture must be rethought, perhaps using the collected capabilities of all endpoints collaborating together. This collaborative approach can reduce the attack surface of each endpoint by using micro-virtualization, continuous monitoring, and execution correlation across all endpoints.

EA: What do you mean by micro-virtualization? How does that work?

SC: Micro-virtualization is a security technique that uses CPU-enforced isolation using CPU features for virtualization to invisibly isolate each task that processes untrusted files or sites. This includes each tab in the browser, each mail attachment, downloaded files or files from a USB key, media, and executables. An isolated task is called a micro-VM, and contains no valuable information or credentials, cannot access the enterprise network or SaaS sites, and is discarded when the user closes the task, eliminating malware persistence. The user is unaware of micro-VMs. Security in a micro-VM is further enhanced by monitoring for signs of attack. Using this method, the endpoint protects itself and provides real-time intelligence for each attack with a minimum of false alarms.

EA: You said the endpoints collaborate together? What do you mean by this?

SC: We regard each endpoint as a *sensor* in a distributed breach detection system. The way it works is that the endpoint first monitors its own execution to detect malicious execution and share its intelligence with the security team in real-time to accelerate enterprise-wide response. The monitor is protected using micro-virtualization to prevent it from being disabled by malware. The endpoint also self-remediates to remove malware that has executed in a micro-VM. Endpoints then share their attack insights in real time with the Bromium Enterprise Controller (BEC) which correlates them. The BEC immediately and automatically searches all endpoints for evidence related to the detected attack to help security personnel respond to any East-West movement of an attacker through the network.

EA: How hard is it to design endpoint protection that prevents malware, but also allows users to access the content they need for their jobs?

SC: It requires great care in product design to ensure that the user experience for any application that accesses untrusted content remains unchanged. A key design principle is to minimize complexity. The use of CPU enforced isolation, for example,

is attractive because it doesn't require adding a whole bunch of new protection code, as you find with many operating system and application software vendors trying to protect their products. In addition, every effort has to be taken to ensure that existing endpoint agents will function as they always have, and not every endpoint security vendor has been successful in this regard.

EA: How does virtualization play into enterprise protection? Is this how you create separation between, for example, a browsing session and the real operating system?

SC: Ultimately the value of virtualization is granular isolation, which leads to better security. This has certainly been the case in the enterprise data center, where companies like VMware have articulated the security benefits of micro-segmentation for many years. The use of virtual security for endpoints is also becoming more evident, but there are some technical differences. For example, micro-virtualization on the endpoint is task-specific, with each browser tab or document operating in its own micro-VM and associated environment. In addition, CPU-enforced protection on the endpoint, afforded by virtualization, is key to granular isolation and reducing the attack surface. So it should come as no surprise that operating system vendors like Microsoft are beginning to recognize that virtualization is a building block of a more secure operating system in the future.

EA: Is it reasonable to say that Microsoft PCs are less secure than Macs? And a follow-up question is whether strong task isolation evens the score.

SC: Apple users *think* they are more secure, but the truth is that Macs represent such a small fraction of *enterprise* endpoints that they really aren't particularly interesting to most attackers looking for valuable business assets. Admittedly, Macs are growing more popular in the C-Suite, and since executives are high-value targets, we will see more targeted Mac attacks. But remember that it isn't only the operating system that's the problem. Any vulnerable application is enough for a good hacker, and every operating system has its share of issues. Fortunately micro-virtualization can be used to protect both PCs and Macs – and my belief is that this may be the strongest hope any security team will have to gain full security control of their endpoint PCs. One parenthetical comment is that protection of endpoints running old, unusual, or proprietary operating systems remains a challenge. Since this includes important endpoints such as POS devices and SCADA controllers, the endpoint security industry will have to continue improving and innovating in the coming years.

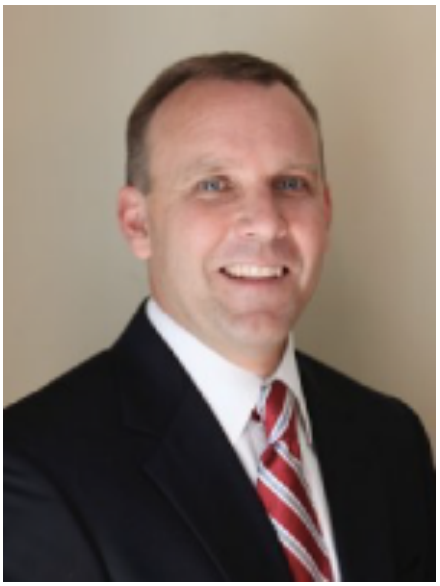

Virtual Software-Defined Segmentation

Mitigating Weaknesses in Modern Enterprise Security Perimeters

David Keasey, CEO of Catbird

As enterprise teams have come to recognize that the traditional perimeter no longer works at stopping or even slowing advanced persistent threats from nation state actors, the need for alternative solutions has never been more urgent. Since 2013, David Keasey has focused his efforts on helping enterprise CISO teams build software defined segmentation solutions, which are designed to support advanced protections for emerging hybrid cloud infrastructure. Techniques that make use of the power of virtualization are likely to be compatible with clear trends in cloud-based data centers, as well as software defined wide area and mobile networks.

EA: Is virtualization truly a reality in the modern enterprise? Or are most businesses still operating in the more traditional manner?

DK: We definitely see accelerations in virtualization and hybrid IT, which involves combinations of private cloud workloads across multiple hypervisors, public cloud workloads, and bare metal. While enterprises continue to have significant legacy workloads on bare metal in traditional data center environments, those are static, and new applications are being deployed using a different model. Our team at Catbird has met with some of the largest financial institutions, retailers, government agencies, and global leaders in other industry verticals over the past year, and nearly every one of them has IT projects on-going that are focused on building next generation cloud infrastructures for their companies. Nearly every one of them is a multi-hypervisor, multi-cloud project where they would like a single pane of glass to manage automated deployment, monitoring, and enforcement across all workloads in a consistent manner. Most of these next-generation cloud projects are also looking to deploy a micro-segmentation strategy across all workloads.

EA: Do you see any differences in the adoption rates? For example, are large companies moving to virtualize more quickly – or perhaps less quickly – than their smaller counterparts?

DK: I think large companies have a greater ability to build large project teams to evaluate options than smaller companies. We've had meetings with global financial firms, for example, where literally a dozen or so business units might be represented in the meeting. This is promising, because it demonstrates commitment to the virtualization approach. But what we *have* seen sometimes, unfortunately, is that because virtualization represents such a new approach, it can lead to challenges in the alignment of goals and the prioritization of requirements amongst the various teams. This can really hurt these organizations in moving their projects forward, so the security industry needs to provide them with improved support, platforms, and processes – not to mention education and training on cloud and virtualization.

EA: Does software defined networking make things better for security teams? Or do you think that it could introduce serious new security risks?

DK: SDN has the potential to simplify and improve the application of security policy. I would argue that the foundation of every SDN provider's pitch is security, and in the end, I do believe the technology will have the ability to deliver stronger security, so long as the customer adopts the security components of the SDN solution, many of which can also be delivered without implementing SDN. Because of the prevalence of legacy systems, and the relative immaturity of many teams to understand and adopt SDN, there are a lot of partial implementations out there today. For example, VMware has been successful in deploying its NSX capability to many customers, but many of those customers are probably still struggling to properly implement the NSX Service Composer and its partners' security controls. And without doing so, these customers are not really enhancing security and receiving the benefits of a micro-segmentation strategy. Companies like Catbird and others in the software-defined segmentation space can deliver the value and enhanced security of micro-segmentation without requiring full SDN adoption. And for customers who truly want to take full advantages of the agility and efficiency of SDN in their networking strategy, then the synergy of improved security and SDN benefit can be achieved.

EA: Just about everyone is talking about East-West traffic threats. What's been your experience working with customers of the Catbird solution? Is it really possible to virtually segment cloud workloads?

DK: Yes. It is absolutely possible, and even straightforward, to logically segment virtual workloads to achieve stronger security overall, and to provide visibility into East-West traffic, which is not typically visible with perimeter based security – and this is an important goal. Consider, for example, that East-West traffic represents over 80 percent of all network traffic. Furthermore, data breach reports from

Symantec, Verizon, Mandiant, and others provide similar statistics regarding the number of days bad actors are inside a network before being detected. That could be as many as two hundred or so days for a breach to be lurking in an environment where most of the traffic is East-West. In my opinion, implementing a software-defined segmentation solution like Catbird provides the necessary visibility, monitoring, and enforcement to really make an impact on this problem by leading to earlier detection of anomalous network traffic.

EA: Catbird supports so-called TrustZones for virtual grouping. Can this grouping be extended across the enterprise to non-virtualized assets?

DK: We are working on two solutions to extend our TrustZone concept across the enterprise, not only to non-virtualized assets, but to deliver on our vision of a consistent approach to automated security policy deployment, monitoring, and enforcement for any workload on any platform. And this includes hybrid IT infrastructure such as private clouds, public clouds, bare metal, and containers. Today, we can extend our TrustZones to adjacent physical assets, which include bare metal workloads connected with the VMware or OpenStack clouds where Catbird is deployed. We've also done work with leading SDN providers, which allows us to extend to all physical assets via the SDN controller. A huge advantage of Catbird is our ability to deliver the functionality we have today in a non-intrusive agentless approach. That said, we believe an agent will be required at some point to cover all bare metal not covered by SDN and some public cloud platforms. At the end of the day, I think enterprises are desperate to find "one pane of glass" through which they can achieve a consistent approach to security across the enterprise.

EA: Do you see many differences in the required security solutions for proprietary cloud operating systems like VMware versus emerging open source virtual platforms?

DK: From a Catbird perspective, we believe enterprises should aim for a consistent approach to automated security policy deployment, monitoring and enforcement across all hybrid IT platforms, including all hypervisor variations. We are working diligently to make this possible in terms of improved visibility into the workloads and automated protection of the workloads. Yes, there are differences in how Catbird integrates with the various platforms an enterprise may elect to adopt, but we strongly believe enterprises will ultimately choose the solutions provider who can deliver the consistency of how policies are defined and applied across all platforms, ultimately managed through a single pane of glass.

EA: You've been in the cybersecurity industry for many years. Are you optimistic that the global IT community can start to more effectively counter these nation-state APT attacks? Or do you think we might simply be doomed?

DK: I am optimistic, but improved cyber security is dependent upon enterprises embracing the organizational change in conjunction with tools adoption to be more

effective. The status quo enterprise perimeter approach is demonstrably inadequate at stopping cyber attacks and new infrastructure protection platforms and tools such as from Catbird now exist which can significantly improve the enterprise cyber security posture with a full defense in depth approach. We're looking forward to seeing organizations lean into the new paradigm and improve both security and efficiency.

Continuous Protection for Cloud Infrastructure

Offering Advanced Security Services for Cloud Workloads in Agile Dev/Ops Environments

Carson Sweet, CTO of CloudPassage

Moving development focus from traditional waterfall to Agile Dev/Ops introduces the need for corresponding shifts in cyber security. Similarly, moving enterprise applications and systems from on-premise hosting to virtual workloads hosted off-premise in the cloud also requires changes in security. The best available solutions to dealing with these shifts offer means for using automation to provide continuous protection through the entire cloud workload lifecycle. These solutions also focus on improving visibility for developers, administrators, and users into the protection aspects of virtual systems. Obviously, this helps immeasurably with compliance and other required security tasks.

EA: How important is it for new and existing enterprise cyber security solutions to support Agile IT?

CS: It is *very, very* important. Agile IT can be viewed a broad, umbrella term for a combination of agile application development, virtualized cloud infrastructure, SaaS, DevOps, continuous integration and continuous delivery, and on-demand, as-a-service delivery. This combination of technologies and operational constructs has created a fundamental shift in how information technology is delivered to enterprise and consumer users. In fact, in my opinion, this the largest technological shift since the IT community moved from mainframe to client-server. When a shift this significant builds up momentum, security is often a top concern – so it's critical for cyber security solutions to be "agile-savvy." Those who don't pay attention to this trend are soon going to find that security architects pass them over or relegate them to legacy environments.

EA: Since the essence of Agile development is speed, do cyber security vendors have to adjust their concept of time? Specifically, is "continuous" the new normal?"

CS: Yes, I do think that continuous is the new normal. In fact, that is a great way to think about it. When CloudPassage started in 2009, we watched the way that agile application development completely changed the pace of competition. And now, compute, storage, networking, and now security, all share this pace, meaning that in order to keep up competitively, enterprises must get agile. Based on our work with dozens of the companies in the Fortune 1000, we've learned that they also know this.

EA: What's been your experience with enterprise teams moving their applications and systems to cloud? Do any specific trends or observations come to mind?

CS: The biggest unspoken trend I see is that every enterprise goes through a phase where they aspire to agile, but as with most technology trends, they quickly realize that there's no magic wand. Three years ago, every large enterprise was convinced they'd be able to do it all themselves. Over time, however, they realized that a system administration is not a Dev/Ops engineer, and that a traditional, vertically scaled application can't be forklifted to a cloud environment. They also learned that cloud infrastructure requires a skill set that's in low supply and high demand. As a result of these reality checks, we've seen a big uptick in public IaaS adoption with large enterprises. This is interesting considering that enterprises consider agility critical enough that they must immediately adopt an alternative approach to getting there.

EA: Every cyber security expert likes to complain about that subset of compliance managers and regulators who don't understand technology. How in the world will these individuals ever come to understand and approve the use of complex structures like dynamic micro-segments on virtual cloud workloads?

CS: Let's hope that this situation improves with the passage of time and the installation of new leadership. Actually, the challenge with auditors and regulators catching up to technological change is really no different in this IT delivery evolution than in previous situations. Or instance, in the late 80's and early 90's, compliance teams struggled with the distressing idea that RACF and Top Secret were no longer the centers of the security universe. They soon got caught up with IP and Internet technologies by watching the industry leaders, who in turn helped drive some of the earliest compliance standards around that space. So it will take industry leadership to define what's possible and tenable, along with sufficient time for this knowledge to percolate throughout the industry. History doesn't always exactly repeat, but I think it does echo – and if that's the case, we can probably expect financial services and telecommunications carriers to be the early leaders.

EA: Do you think that virtualization will really help the enterprise thwart cyber threats? We all know it saves money, but will it save assets from being attacked as well?

CS: It cuts both ways. One big benefit of virtualization – and the related, but more specific technologies of containerization and micro-services architecture – is that the infrastructure can become a moving target. This of course depends on the deployment model, because vanilla virtualization is usually more of a moving target than bare metal, but not as much of a moving target as containerized micro-services. Another benefit is that virtualization and the related infrastructure orchestration can drive new heights of consistency in deployment and configuration. However, this is an example of how something can cut both ways. That is, a bad configuration decision or mistake can turn a virtual machine template into the equivalent of Typhoid Mary, spreading exposure rapidly through the enterprise. So as with most new technologies, there are upsides and downsides.

Cloud-Enabled Intelligence for Protection of Endpoints

Next Generation Endpoint Security Combines Intelligence, Cloud, and Support for Analytic Hunting

George Kurtz, CEO of CrowdStrike

The nature of cyber security had shifted from static protections such as early signature-based anti-virus to more proactive, real-time protections based on observed indicators of attack (IOA). Leveraging the cloud and graph database technology is the primary means for powering such real-time security. Combined with machine learning and intelligence, it makes for a potent prevention platform. Next generation endpoint security thus works best at the intersection of cloud, machine learning, and intelligence, offering a solution that supports real-time proactive security as well as after-the-fact hunting for indicators should a breach occur.

EA: What do you see as the definition of an endpoint today? It used to be simple, namely your PC or laptop. But today, does the definition include mobiles, IoT devices, industrial control elements, cloud workloads, and on and on?

GK: We do have enterprise customers who maintain a conventional view of endpoint definitions, where employees are mostly using Windows PC connected to perimeter-protected LANs running enterprise software tools such as the Office suite and Active Directory. In these environments, the currently evolving notion of the endpoint, including industrial control or IoT devices is interesting, but not relevant to their day-to-day world. This is true for mobile as well. Certainly, we all have our iPhones and Android devices – and I guess Blackberry too, but these are often not viewed today as being true endpoints in many companies and government agencies. At CrowdStrike, we know this will change, so we have the obligation to our customers to both serve their existing needs, and also to be helpful and ready to support evolution to a broader context. Clearly, employees of companies of all sizes

will begin viewing their tablet on par with their laptop or PC, and our technology is designed to support such evolution.

EA: What is the specific role of machine learning and graph analytics in protecting endpoints? Is this the secret to replacing signatures with something more acceptable?

GK: There is certainly nothing inherently evil about signatures. As you know, for many years, the only technique we had in the protection of endpoints was the development of signatures based on deep analysis of viruses, worms, and other malware. With the development of variants, however, the work equation shifted and it became significantly easier to develop a modified version of some malware than to develop signatures. This work gap between the offense and defense required modification in the protection strategy – hence we began to use machine learning over our ThreatGraph (which handles 20B events per day of threat data) rather than signatures into our security tools. This really became a major help, because it opened all sorts of new possibilities, including the potential to stop malware at both static and runtime. More importantly, we can look at the attack kill chain and identify behaviors that do not use malware. In fact, many of the most virulent attacks don't use malware, but instead use things like credential theft and social engineering. No AV is going to detect this activity. We like to say this: "Stop the breach and go beyond just stopping malware."

EA: If a nation state figures out an APT method that is detected and reverse engineered, can they continue to make straightforward adjustments in the attack to hide from security teams? Or does the security and behavioral analytic process prevent this from working?

GK: Well, you can't prevent nation states from figuring out and doing whatever they might decide to do. They are capable adversaries, and will, unfortunately, just get even better. And yes, they will try to develop slight variants, simply because that is so easy, and is the currently accepted offensive methodology. But our approach at CrowdStrike is to develop static and dynamic technologies, based on threat intelligence and deep understanding of both the endpoint and the operating environment, that stop broad classes of attacks, without having to pick apart the malware for a specific, easily changed identifier such as a file name. Our Falcon platform drives this type of solution to the endpoint with the goal of continuous visibility, which is a big change from early anti-virus. And of course, our threat solution is delivered via the cloud, which really does change the game, given its ubiquitous nature. Finally, our teams of expert adversary hunters are essentially watching our customers' backs on a 24/7 basis, which has proven to be invaluable in dealing with a capable adversary who changes tactics. Put all this together, and the defensive solutions to endpoint attacks are so much better than in the previous generation.

EA: You are one of the pioneers in anti-malware techniques. What are some of the trends you see in this aspect of cyber security technology?

GK: Not only are we pioneers of anti-malware technology, but we are the only company to unify next-gen AV endpoint detection and response (EDR) and managed threat hunting into a native cloud platform with only one agent. Yes, that is right – just one! Moreover, the new focus on cyber security hunting is a welcome trend. The image of the hunter is exactly the right one for modern enterprise cyber security. Armed with excellent tools, and we think our platform is among the best, the cyber hunter from our Falcon Overwatch team is both doing preventive work, looking for early indicators, as well as response work, and looking for indicators of on-going or previously initiated attacks. It's interesting that hunting combines both active and passive, preventive, and responsive, and also automated and human defenses into a common, integrated solution. At CrowdStrike, we not only support the enterprise hunter, but we backup our solution with some of the best expert adversary hunters in the business. Now, granted, a smaller organization will not have the budget or desire to hire a bunch of cyber hunters. This is obvious, and our platform and endpoint solution is designed to provide them with world class protection offering proactive support that also leverages crowd sourced knowledge to make up that gap.

EA: Do you think we'll ever see a true Cyber 9/11? Is this inevitable given the asymmetry of the cyber security equation?

We have to be careful about how we characterize cyber risk. As much as it is a critical threat, there is no loss of life involved and we don't like to use hardline analogies like 9/11. You're correct about the asymmetry. We all know that it's easier to attack than to prevent. At the core of today's advanced defensive approaches is the belief that there is no silver bullet. New techniques like behavioral blocking, machine learning, continuous monitoring and proactive threat hunting make it significantly harder for the attacker to succeed. If we mature our defenses to match or surpass the maturity of the attacker, then we stand a much better chance of preventing mega breaches. In short, we all have to raise our game. At CrowdStrike we deliver the people, process, technology and intelligence to make it easier for organizations to do just that. We also leverage advanced techniques like machine learning and threat graph technologies that amplify security resources and scale capabilities to drive defenses that are faster, more agile, and ultimately more successful than the offense.

Secure Access in a World Without Network Borders

Initiating Secure Connectivity Across Network, Virtual, and Cloud Boundaries

Barry Field, CEO of Cryptzone

Secure access to computing resources is one of the core goals of computer security. In modern cyber security, it remains one of the great challenges, especially as the traditional enterprise perimeter has dissolved into virtual and cloud infrastructure. Creating a proper level of secure access requires technologies such as strong authentication, authorization, encryption, and fine-grained access control. Adding to the challenge, these controls must be provided over any type of network in a variety of different settings. Implementing these controls can be difficult, but the results are worth the trouble: Reduced complexity, more granular controls, and higher levels of user satisfaction result from secure connectivity across networks with virtual borders.

EA: You've referred in the past to secure access as requiring a 'segment of one.' What do you mean by this?

BF: Traditional network security tools typically allow end-users more access to the network than they actually need. In addition to allowing a user to see more than necessary, traditional access tools do not always limit users to what they can and cannot do. Secure access in today's environments instead requires security platforms and tools that can check a user based on their context. This includes identity, device information, user location, network being used, and application sensitivity, driven by dynamic and easily configured policies. By limiting a user to the assets required to do his or her job, a so-called 'segment of one' is created. This ensures that businesses can control network access at a more fine-grained level, limiting users to only authorized resources, whether on-premises or in the cloud, and rendering everything else on the network invisible.

EA: Is there a way to check the integrity and security characteristics of an end-user device before it is permitted access to some resource?

BF: Yes, that is typically referred to as device context. Advanced authorization methods should include the ability to capture the posture and context of each session. For example, before allowing access, an enterprise may want to check whether the device has anti-virus software installed, what time of day it is, what the location of the device is, and other variables. These types of requirements are more and more often found in enterprise security policies.

EA: In your opinion, do you think enterprise CISO teams should be comfortable accessing public cloud assets for sensitive or critical business applications?

BF: We've seen tremendous uptake with organizations of all types and in all sectors moving applications to Amazon Web Services (AWS), Microsoft Azure, and other public cloud environments. The business value of using public cloud services is proven, and to not use it because of security concerns is a mistake. Using 'segment of one' granular access controls, such as from the Cryptzone team, businesses can use dynamic controls to secure their select IaaS environments.

EA: Enterprise CISOs today usually have to accept the existing architecture of their organization – and this is not always such a great set-up. If they want to construct a more secure virtual network on top of their existing network, what are the steps they must follow?

BF: CISOs need to consider using technologies that enable them to use the same access policies, whether they are trying to control third-party or employee access to a network-based application or one that's in the AWS or other public cloud environment. These technologies should seamlessly control access on a per-user-session basis to significantly reduce the attack surface and enable secure. They should also ensure controlled access to privileged users and third-party organizations, regardless of whether those resources are located on-premises or in the cloud. These considerations should greatly assist the CISO team trying to construct a more secure virtual network.

EA: Most companies have come to accept stronger authentication as a requirement. Do you see one-time password use growing – and is the mobile device the preferred device?

BF: The continued widespread adoption of mobile devices will erode the use of on-premises PC. And yes, we see one-time password use growing, as well as other multi-factor authentication choices, based on risk-driven policy decisions that you should be able to adjust on the fly. For instance, if an employee is trying to access a low-risk application hosted on-premise and that they access at about the same time every day as a typical activity for their role, you may not want to invoke a multi-

factor authentication sequence. If however, that same employee appears to be accessing a more sensitive application at an abnormal time of day and from a different location, then by all means, the security teams should dynamically invoke one-time passwords or another multi-factor authentication method. Network security relies on dynamic controls to ensure secure access to the resources within your network. Utilizing solutions that provide this help will always reduce the potential damage that a determined cyber adversary may cause in the event of a breach.

Advanced Methods for Predicting Network Threats

Combining Analytics, Threat Intelligence, and Machine Learning Into World-Class Anomaly Detection

Tom Caldwell, Executive Vice President of CyberFlow Analytics

The disappointing results found by most enterprise security teams with traditional intrusion detection and intrusion prevention systems can be attributed largely to the simplicity with which clever adversaries can bypass familiar signature patterns. But when the defender uses advanced methods for building behavioral maps, developing machine learning, and combining different network security profiles together into a comprehensive network security prediction and detection model, it truly becomes more difficult for even the most determined adversary to successful attack an enterprise.

EA: What is the technical difference between predictive analytics and reactive response?

TC: Many enterprise security teams now realize that they are not properly prepared for the so-called "unknowns of the unknowns" with current cyber security tools such as SIEMs. A reactive response involves focusing on *not losing* the cyber battle and hoping that legacy security tools will provide sufficient protection. This approach usually includes waiting until an attack actually occurs so that another rule can be added to the security tools and the SIEM. In contrast, predictive analytics are designed as an early warning system to detect anomalous activity in the very early stages of the kill chain. The goal should therefore be to focus on proactive means for detecting and stopping breach activity before a crime has taken place. It's better to catch the bad guys early in the process while they are just snooping around or trying to set up shop to plan out a long range APT attack.

EA: What do you see as the role of threat intelligence in optimizing the analytics used to detect cyber threats?

TC: Threat intelligence is a critical element to optimize analytics. But security teams must understand that just because something is odd or anomalous does not mean it is dangerous and adversarial. Instead, you need *context* to understand the risk of an anomalous behavior. Threat intelligence can help rapidly identify external sites that are connected across a network to the internal anomalous communication. The process of connecting the dots is important as you map out adversarial activity across your internal and external network. Threat intelligence can also be used to help identify families of behaviors using supervised machine learning techniques. For example, a first stage of unsupervised machine learning can separate out normal from abnormal behaviors on your network and devices. A second stage of supervised learning can then help detect similarities in behaviors from morphed Malware, RAT, or Ransomware behaviors. The end result is detection of strange and adversarial behaviors by monitoring the network with predictive analytics.

EA: What is the role of machine learning in cyber attack detection? Is this an essential element of future preventive defenses?

TC: Machine learning becomes more critical as more devices are connected to the network. A human being cannot begin to understand the normal behaviors of all assets on a network and the ports they normally communicate with. The first type of machine learning is known as unsupervised machine learning, and it is used when you don't know "bad," but you know "different." This approach is critical in detecting anomalous and unknown adversarial behaviors on the network. A second type is supervised machine learning, which is used when you know a behavior is likely to be "bad" and you want to further refine what type of adversarial bad behavior it appears to be associated it. The rules used in a SIEM tend to be less effective because they look for exact behaviors in the network. Machine learning loves Big Data and becomes more accurate with the more data you feed it. Machine learning analytics can look for similarities in behaviors as they become better at classifying good from various types of bad. SIEMs typically cannot handle Big Data well.

EA: How hard is it to combine different behavioral profiles into a more holistic view of system behavior? Does this require complex mathematics?

TC: The process of defining different viewpoints using a variety of behavioral models is something we call "fusion." You can have machine-learning models for understanding the behavior of protocols on the network, and this involves looking at internal-to-internal communication differently than internal-to-external communication. You might also want a model that looks at client-server-by-port behaviors, or a model that can study one device and the ports it normally communicates over to catch a workstation being turned into a server such as an

SMTP relay. All of these various models provide opinions on your network, devices and users, thus helping to highlight the difference between normal behaviors and potentially adversarial behaviors. The process of fusing these opinions together into a security alert typically takes a second stage of analytics and threat intelligence processing to weed out the positive anomalies from the neutral or negative anomalies. Negative anomalies could be the result from operational problems or adversarial activities. The goal is to quickly assess risk on the negative anomalies and then turn it into actionable intelligence. For industrial IoT environments, operational risk can be as dangerous as adversarial risk.

EA: Many cyber security teams are now focusing on a technique known as hunting. Any thoughts on whether this is the right approach?

TC: Hunting groups in security operations centers or in vendor security research teams take a more personal view and go after the attackers themselves. The intelligence they collect can help organizations become better prepared for the advanced techniques these attackers may commonly use in their missions. However, these hunting teams are typically staffed with highly advanced security analysts. This caliber of security professional is hard to find, hard to hire, and in great shortage of demand.

EA: What general trends do you see on both the offensive and defensive sides of the cyber security equation?

TC: I see a need for automation and intelligence. There is a shortage of trained security staff, and there are never enough hours for current security staff to handle all the potential security issues that need researching. What is needed is an industry transformation of software that enables automation of security tasks with embedded intelligence to perform many of the manual tasks of a security analyst. As this intelligent automation software for security becomes integrated into the security framework, the APIs take on an essential role to ensure that the entire security infrastructure works properly. We saw this transformation happen in business and operations support systems (OSS). Now we need a transformation in automating integrated offensive and defensive processes of the cyber security equation.

Fighting Malware with AI and Machine Learning

Continually Evolving the Best Algorithmic and Computing Techniques to Fight Malware

Stuart McClure, CEO of Cylance

For many years now, artificial intelligence and machine learning have consistently topped the list of areas in which pundits have predicted advances in computing. In some sense, the promise of these technologies is irresistible, offering super-human means for performing computation and reasoning that was previously unthinkable. And if you doubt that such advances are possible, just tap in a Google search for instantaneous access to some piece of information that two decades ago would have required months of research to obtain. So now, the promise of AI and machine learning has come to cyber security, and the big question now is whether these techniques will give a real advantage to the defense.

EA: Can we say that traditional anti-virus software with its base of signature patterns is officially dead?

SM: We can certainly say that it's actively failing customers around the world every minute of every day. The reason that the entire endpoint detection and response market grew up was to try to help security teams find what last-generation anti-virus keeps missing and clean it up as fast as possible. The problem is that these are arson investigation tools, and there's no time to do an investigation like that when the house is on fire. That's why prevention is so critical.

EA: How do you explain to a non-expert how artificial intelligence can be applied to problems of cyber security?

SM: Information security today is a numbers game. Artificial intelligence and the mathematical algorithms we're building are ideal for scaling to meet the number of attempted attacks that organizations are experiencing every day. The headlines will tell you that companies locked into older endpoint security technologies are losing the battle, because the number of threat actors is increasing, the number of malware mutations and attack types is increasing exponentially, and the number of attempted attacks on increasing numbers of devices is also growing. People don't scale, and budgets will never be infinite. So if you're relying on either your team of experts or the signature-coding teams inside traditional anti-virus companies to scale to meet the crushing number of attacks, then you can't win. Machine learning makes it possible for Cylance technology to predict whether something is an attack based on hundreds of thousands of properties learned from earlier attacks. Current industry-standard techniques, such as signatures, heuristics, and behavior monitoring rely on simplistic, easily evaded data points. Best of all, algorithms can convict a file as bad or good in milliseconds, they don't need coffee breaks, and they never have a bad day as humans sometimes do.

EA: Regarding machine learning, isn't that really just traditional expert system rule-based construction? Or is there really something new here?

SM: There really is something new, and that's why our application of artificial intelligence to anti-virus is able to operate disconnected from networks and the cloud, and is even protecting air-gapped networks inside critical infrastructure installations. Traditional anti-virus products – in fact, all other anti-malware products – can allow a zero-day in and gut your road warriors' laptops within about a day of being disconnected from the cloud or corporate network to get the daily update. In fact, the advancements in the last decade in machine learning, including ours at Cylance by our team of data scientists and mathematicians, have fundamentally changed the way we interact with technology. These genuine breakthroughs that have taken machine learning to new levels are responsible for our 99+% conviction rate when encountering brand-new malware – malware that might have been created only the day before. Our Chief Data Scientist Matt Wolff came from the NSA's TAO unit and often shares our breakthroughs in areas such as deep learning on disassembly data and new applications for data exfiltration at security conferences such as Black Hat. What our team discovers each week gets used for continual evolution of our math model and algorithms for use on the endpoint.

EA: Empirically, are you seeing better results in the enterprise regarding malware in critical systems?

SM: If by "better results" you mean "better malware protection in the enterprise on critical systems," then yes. But generally, this is only true because they have adopted very heavy-handed operational controls of those systems. The change control process requires multiple layers of approvals and still misses major attack vectors

such as stolen private certificates as well as malicious insiders posing as programmers. Our technology at Cylance is designed to help enterprise users truly achieve better results without having to change business processes or adopt heavy controls.

EA: You've been in the cyber security industry as a leader for many years. What are the key trends that you see occurring now and into the next few years?

SM: We're definitely beginning to see the rise of ransomware as a service. It's a simple matter now for any random threat actor to hire a mercenary hacker to buy some ransomware off the shelf and make a minor modification to turn it into something that can take a business offline for days or weeks. The worst of it is that federal law enforcement are saying these types of attacks can almost become "the perfect crime" – namely, untraceable when the criminals are operating through third parties and being paid in Bitcoin.

EA: Do you ever see the defense catching up to the offense in cyber?

SM: Yes, I think we will finally see, for the first time ever, the defense catching up. And this can be best achieved by applying AI to the prevention of threats on the endpoint. Leveraging Cylance's technology, we can foil 99+% of cyber attacks without ever allowing the malware to run. So the defense *can* catch up and continue to harden their security posture, resting assured that the ultimate mission of nearly every one of the tens of thousands of attempted cyber attacks will be thwarted.

Data-Centric Visibility To Prevent Cyber Attacks

Shifting Focus from Perimeter to Data-Centric Controls Including DLP Offers Advanced Threat Protection

Ken Levine, CEO of Digital Guardian

The image of walls coming down around enterprise networks leads to the obvious conclusion that CISO teams must focus on data. The challenge is that visibility into structured and unstructured data in an enterprise requires a fundamentally different paradigm than the existing focus on servers, systems, and networks. Data loss prevention (DLP) is a start, but the modern enterprise security teams needs to have tools and processes that ensure full visibility, accountability, and protection from cyber attacks for all relevant business information.

EA: It is obvious across the cyber security community that DLP projects are one of the most important priorities in the enterprise. What do you think are the main factors driving this focus and interest in DLP?

KL: We see three main factors in the growing acceptance that DLP is a critical control in enterprise data protection. First, we see a growing acceptance that systems have already been breached. This recognition is certainly not good news, but it does produce the healthy view that with malware already existent in the enterprise, that something had better also be in place to prevent sensitive data from leaking out – which is where DLP solutions are well-suited. Second, we see a daily drumbeat of enterprise attacks, many of which have involved prominent companies losing a lot of business data, sensitive information, and other assets. The Sony attack, for example, shone new light on the business impact of a data breach, and the C-suite seemed to notice. At Digital Guardian, we saw a noticeable uptick in inbound interest after the Sony attack, because security teams began to recognize how important it is to avoid leakage from endpoints, systems, and networks. Finally, a third factor is the improved delivery of DLP products and services to enterprise customers. For example, at Digital Guardian, we began offering an outsourcing DLP

solution, because we saw the scarcity of security talent to manage the technology and the ongoing care and feeding required to have a successful data security program. Solutions like this are designed specifically to make it easier to deploy and use DLP.

EA: How exactly is cloud DLP different from on-premise data loss prevention products and services?

KL: Both cloud and on-premise DLP have the same goal, and that is to prevent data from leaking into unauthorized hands, either as a result of unintentional or deliberate action. Cloud-based DLP solutions are focused on protecting customer sensitive data as it moves to and from the cloud. This complements similar offerings focused on providing DLP for the network, enterprise, and endpoint. These are all offered either on-premises or outsourced to Digital Guardian experts. Our product, Digital Guardian for Cloud Data Loss Prevention, integrates with leading cloud storage providers such as Box, Citrix and Microsoft to extend DLP policies to the cloud. This solution is a good example of the type of capabilities security teams should be looking for in their DLP. For example, they should be demanding accurate sensitive data discovery for cloud storage; they should demand continuous protection of files that have been uploaded from the cloud; they should require automatic remediation according to enterprise policies; and they should demand instant alerts sent to the appropriate administrator and data owner when some event or action require attention. As more corporations rely on cloud technologies, it's important they take the proper steps to protect their sensitive data as it moves outside the traditional IT security perimeter.

EA: What are some of the biggest challenges and opportunities you expect to see in the cyber security market over the next five years?

KL: In the next five years there will be an increase in sophisticated targeted attacks, which will force a convergence to happen within endpoint detection and response, endpoint protection platforms, and endpoint DLP. This might be driven at the product level by mergers and acquisitions, as the larger vendors integrate the capabilities of smaller, feature-specific security vendors. Another trend is that given the continued talent shortage, security teams will come to rely even more on automated protection platforms. CISO teams can no longer employ disparate systems that require scarce resources to log into and manage a variety of separate panes of glass. They will instead demand an endpoint command center where they can access endpoint health, and instantly react based on their diagnosis – whether that means applying a browser patch, fully quarantining a device to suspected malware infection, or some other actions. A third trend will be greater data awareness at the endpoint. Security professionals will have complete visibility into the sensitive data accessed on the device and enable policies and controls to protect it at rest, in motion and in use.

Raising the Security Bar on User Authentication

Advancing the goal of trusted access for the enterprise move to the cloud.

Dug Song, CEO of Duo Security

The advantages of passwords are well known by now: Simplicity, portability, low cost. But the reality is that modern cloud and enterprise applications are so highly vulnerable to malicious attack that two-factor authentication, a second layer of security, has become a *de facto* requirement. Yet, so many existing on-line, Internet, mobile, cloud, and enterprise systems and services continue to neglect this fundamental control, despite compliance regulations put into place to enforce its use such as PCI, HIPAA and FISMA. Without stronger forms of authentication, experts agree that continued break-ins, leakage, and even destructive attacks are likely to occur.

EA: Do you still use passwords for any of the services you personally use on a daily basis?

DS: Yes, but with some modern coping mechanisms. As you know, every service still defaults to password-based logins, and this is a problem because passwords must be hard for attackers to guess, but easy for users to remember. This presents a basic compliance defect – namely, something that is hard to follow and impossible to enforce. Nevertheless, there are techniques now that can help make them safer and easier. Password managers, for example, like LastPass or Dashlane help users cope with this complexity, but the burden remains to find and use such tools, which are protected, by the way, with a single master password. A more powerful trend involves the open interoperability standards that reduce the number of logins and simplify identity management. OpenID, for example, enables Websites to federate the login credentials of consumers from their Google, Facebook, Twitter accounts. Similarly, enterprise single-sign on via SAML or OIDC is a common technical

approach; it's implemented on the Duo Security platform. So with these new single sign-on and federated approaches, the days of having separate passwords for every application will quickly fade away. To paraphrase Mark Twain: It's better to put your eggs in one basket, but you have to watch that basket!

EA: How about introducing additional factors? Isn't that really the best way to make user authentication more trusted?

DS: Yes, and the most common approach involves establishing a level of trust in the device being used. Enabling Web applications, for example, to "remember such-and-such computer for some number of days," and to not require password login each time that application is being used, is both convenient for users and better for security. This works best if the trust enablement can be established for devices to properly audited and HTTPS-protected services, and if logins are disallowed from unknown or unsafe devices, so that stolen passwords are useless to an attacker with an untrusted device. Some consumer services already alert on new devices being used to log in to an account, but newer enterprise platforms such as Duo's can actually stop them based on device identity or security profile. And this brings us to the strongest method, which involves two-factor authentication. Wearing a belt and suspenders might be out of style and redundant for your pants, but on the Internet, requiring a second factor of authentication is essential. Increasingly, users understand the need for this from their consumer experiences, such as when they use a credit card at a gas station pump requires the entry of a zip code. Choosing services that at offer two-factor authentication is for me – and I hope is for everyone reading this interview, like choosing restaurants based on the cleanliness of their environment. Two-factor authentication should be viewed as evidence of operational excellence and good hygiene.

EA: This sounds great coming from a security expert. But are you actually seeing greater adoption of two-factor authentication across the spectrum of users and business?

DS: For decades now, two-factor authentication (2FA) was limited to protecting only the most privileged accounts for only the highest-risk applications in large enterprises, banks, hospitals, and governments. It was essentially "security for the one percent." But today, 2FA (also known as two-step verification) has become mainstream. At Duo, well more than half of our thousands of customers are green field buyers who have never deployed 2FA before, and nearly all end up deploying to their entire user population, not just admins. With high-profile public breaches since 2010 affecting so many industries, more buyers are driven by tangible risk, rather than just compliance. Cloud and mobile have led the way for 2FA adoption not only by allowing for rapid deployment and global scale – including the removal hardware – but also by allowing for better user experiences. For instance, Duo can enable users to complete their corporate logins with a tap or fingerprint swipe of their smartphone or smart watch – something unthinkable just a few years ago.

EA: In the enterprise, why do you suppose the regulatory and compliance bodies haven't been more aggressive in demanding stronger authentication?

DS: As the technology landscape sees vendors like Google, Apple, and Microsoft bake more security into their systems, broad security initiatives like 2FA become possible in ways regulators never dreamed of. We should therefore expect to see not just more, but better requirements emerge for strong authentication, driven by these innovations in the market. And this is also enabled by the growing adoption of smart phone usage. For example, back in 2009, only 18% of Americans had a smartphone. Since that time, massive adoption has enabled security in ways we've long been waiting for, including local biometric authentication, hardware root of trust, secure boot, secure crypto co-processing, secure software distribution and signing, hardened operating systems, and remote attestation. Further good news is that mobile devices themselves are becoming more secure. It takes about a million or a federal court order to break into an iPhone today, and it's only going to get harder. The latest revision to NIST's federal electronic authentication standard (NIST 800-63) reflects this, as it does away with prescriptive technologies to focus on the characteristics of authentication that define different levels of strength.

EA: Some pundits continue to claim that two-factor authentication is still too burdensome and will never reach mass adoption. Do you agree?

DS: Tens of millions of Touch ID-enabled iPhone6 users would disagree. And it won't be long before laptop and other device are employing similar controls. The real opportunity in 2FA is to enterprise-enable the new generation of consumer security technologies that have been designed for usability. We've gone from telephony to push notifications as smart devices have been adopted *en masse* by consumers, thus offering a better option for 2FA than tokens. Duo was the first vendor to support new security hardware either in smartphones or standalone, such as U2F tokens from Yubikey. We have always been committed to delivering a future-proof cloud-based authentication service, not just a specific form of authentication, which is how we believe organizations need to keep up with new innovations, and changing user expectations. Our goal is to help customers learn how frictionless a user-facing security control can be. It's not only the largest, or fastest-growing companies in the world that are seeing success with 2FA. We have federal customers for whom full-scale deployment, sometimes performed in a day or so, to large groups of users with modest technical skills will produce, only tiny numbers of help desk tickets.

EA: There has been some stir around adaptive, contextual authentication. What is this exactly?

DS: There are many ways to authenticate, some with higher confidence than others. Banks, for instance, whose job it is to know their customer, were the first to adapt

authentication techniques to situational risk for customers at scale. This could mean a teller calling to complete a wire transfer over a certain amount, or requiring a faxback authorization. These were discrete policies applied to govern the risk of specific threats to the business. But in the hands of IT security vendors, this turned into what Jerry Brady, Global CISO at Morgan Stanley, has called the "mystery meat of authentication" – namely, loads of alpha-weighted authentication criteria merged into a composite risk scoring equation, with a step-up to a stronger authentication method, based on fairly arbitrary thresholds. If it sounds complicated, it is. And you can only imagine how users react to inconsistent login experiences, and admins to non-deterministic access controls. We've learned over the years at Duo Security that if you design security to be adopted, you must aim to make things intuitive, elegant, and understandable. Security should frustrate attackers, not users. Integrating multiple security criteria for access control doesn't need to be mysterious. Most organizations understand how they'd like to enforce access; we just need policy frameworks that make this easily manageable, and implementable at scale. We call our take on this Trusted Access, which involves providing simpler, more powerful access controls that are user, device, and location aware, to govern access to any application.

EA: Do you see a day when we basically no longer must remember passwords and where the systems we use provide intelligent authentication without our having to do much more than just show up?

DS: Yes, the Holy Grail is possible, but we have to be vigilant in how we engineer privacy along the way. Authentication still needs to be different than identification, even if done automatically or we'll end up with surveillance programs that mishandle user identities. The difference between authentication and identification is user consent and intent. Authentication requires users to prove identity in ways that can be jointly managed. For example, I can change my password, and I can choose what devices provide a claim to my identity. But I cannot change my mother's maiden name, the places I've lived, date of birth, or my fingerprints, all of which were stolen by attackers in the massive breach at the Office of Personnel Management, affecting so many Americans. Authentication will only get easier and more convenient over time, but I don't believe we'll ever completely get rid of passwords, or want to - we'll just get down to our favorite one.

Multi-Layered Fraud Protection

Supporting the Detection and Prevention of Electronic Fraud Across Industry Sectors

Ricardo Villadiego, CEO of Easy Solutions

As global businesses have moved to the Internet over the past several decades, the inevitable migration of fraudulent activity has also occurred. Perhaps the salient aspect of the fraudster – electronic or otherwise – is resourcefulness. Fraud typically involves an assortment of clever techniques designed to find weak spots in infrastructure and operations. As a result, point solutions cannot solve the entire problem; instead, a collaborative protection suite is required to counter the potentially devastating effects of Web and application fraud on modern business.

EA: There is an old saying that robbers target banks because that is where the money is. Do you think this is a valid description of where fraudsters target their activity?

RV: I think it's a fair assessment, but it's more than that. Fraudsters are going where there's money to be made, certainly, but they are also looking for relatively easy money. So, financial organizations lacking the most rigorous and cutting edge fraud protection solutions are ideal targets. Take the SWIFT attacks earlier this year. Not two months after cybercriminals used the SWIFT system to steal $81 million from a New York bank, another attack has come to light. In this case, hackers infiltrated SWIFT's financial messaging system and sent a dozen fraudulent wire transfer orders to Wells Fargo Bank, asking that $12 million be transferred from Ecuador's Banco del Austro bank to four different accounts located in Hong Kong, Dubai, and the United States. Had a solution been in place that utilized machine-based learning and advanced anomaly detection and prediction, these attacks could have been prevented. Financial organizations, of course, aren't the only verticals being targeted. Hackers, for example, are increasingly turning to the travel industry.

There's real money to be made in targeting frequent traveler and loyalty programs. It's no wonder that enterprises of all stripes are investing in anti-fraud solutions.

EA: Has the business of detecting and preventing fraud become more difficult with the transition to mobile communications in business?

RV: It is definitely the most complex and sophisticated environment we have seen. But just as fraudsters are working to stay one step ahead of law enforcement, fraud prevention solution providers have been working twice as hard to stay ahead of fraudsters. As mobile financial transactions are proliferating, mobile commerce is increasingly coming under attack. A recent report from Forrester Research Inc. noted that it's critical for today's enterprises to seek out fraud prevention solutions that collect mobile sensor data and integrated scoring profiles. At Easy Solutions, we have several solutions that protect mobile devices and effectively stop even the most sophisticated mobile attack in its tracks.

EA: As you said earlier, everyone just assumes it must be financial services being hit the hardest, but can you elaborate on how other industries been hit?

RV: Fraudsters are going where the money is. And that obviously means financial institutions are coming under attack, to be sure, but retailers are also coming under heavy fire. Take, for example, the number of attacks on big-brand retailers a few years back. This included Target, Home Depot, Neiman Marcus, and Staples, just to name a few. Currently, fraudsters are focusing their attention beyond direct financial targets. Personal data is where the money is, and it's especially lucrative for cyber thieves when that data is sold on the Dark Web. Everything from Anthem's medical data breach to the Office of Personnel Management, whose data breach last year involved the personal data of 21.5 million people, is fair game.

EA: Tell me about how fraud takedowns occur. Does this require a lot of analysis by fraud intelligence teams?

RV: We first start the process by proactively looking for fraudulent Websites that might be trying to imitate legitimate sites, often with some slight misspelling variant on a common business or domain name. We also focus on searching for evidence of Websites that are associated with phishing or malware, usually requiring, as you suggest in your question, quite a bit of in-depth analysis by our fraud intelligence team. From there, we work with third parties including ISPs and hosting companies, to actually *deactivate* these malicious sites. Our Detect Monitoring Service shuts down thousands of attacks daily with a 76 percent proactivity rate, meaning that an attack was stopped before our clients or their customers even knew it existed. And with an average take down time of 3.6 hours, we're seven times faster than the industry average.

EA: What advice do you have for businesses that might have some of the tools required for fraud prevention and detection, but that might not possess a full suite of what is truly necessary?

RV: The concept that "one-size doesn't fit all" holds especially true when looking at fraud prevention and detection. Just as no two enterprises are exactly alike, no two anti-fraud solutions are identical. The smart business will look for companies that offer defense in depth strategies with components that can be cherry-picked to suit their individual needs and that will work in harmony with one another. Total Fraud Protection from Easy Solutions is a prime example of this and offers organizations flexible and comprehensive multi-layered fraud protection across all devices, channels and clouds. Even with limited resources, businesses should invest in a solution that monitors brand usage and prevents malicious activity that can lead to account takeovers. They should also be looking for solutions that provide real-time transaction risk monitoring and assessment. Enterprises need to embrace a multi-layered approach to fraud prevention, but more than that, they need to embrace the idea that it's not a matter of *whether* they will be breached, but rather *when*. It's therefore vital that they educate their staff on phishing and spear-phishing, as well as the dangers of opening attachments and BYOD security. Beyond the prevention aspect, however, businesses need to have incident response mechanisms in place for when the inevitable happens, because it will.

EA: In addition to the usual phishing, ransomware or malware attacks, what other types of cyber attacks should enterprise security teams be on the lookout for?

RV: Financial-related fraud makes the headlines, but for any company today, a greater issue might reside in reputation-related threats such as brand impersonations, malvertising, email spoofing, cousin domain registration, and brand infringements. The issue here isn't so much the immediate effects, but rather the long-term damage, not just to reputation, but to the loss of public trust and the ensuing loss of customers. In our opinion, it's just as important for organizations to protect brand value as monetary value. Companies need to monitor for similarly named domains and become hyper vigilant about their social media presence, and this requires looking beyond just Facebook, Twitter and LinkedIn, toward also incorporating Instagram, Snapchat, Tumblr, blogs, and the like. Smart companies also need to protect their email channels from impersonation by employing systems such as DMARC, so that teams can remove those threats before they reach their customers.

Advanced Enterprise Security Platform

Supporting Security Solutions From the Endpoint to the Perimeter and Beyond

Ken Xie, Founder, Chairman & CEO of Fortinet

Just about every cyber security vendor refers to its solution as a *platform*. This is not surprising when one considers the obvious scaling, support, and extensibility advantages of platforms over point solutions. But the term is overused, and most point solutions available today are not really platforms. For example, few vendors take the time to ensure horizontal integration of their solution with other related functions required in the enterprise. Similarly, few vendors take the time to ensure vertical integration of their solution across the architectural evolution, from perimeter to cloud. But when a security platform does provide support for such evolution and expansion, the protection results are powerful.

EA: How important is it for enterprise security teams to know the relative difference between hardware appliances and their virtual equivalent?

KX: Physical hardware and virtual software appliances obviously differ in performance capabilities. Their security features and functions, however, should not. In architecting their networks, enterprise security teams need to consider where it makes the most sense to deploy hardware appliances versus virtual devices, and more importantly, the critical need of having a unified fabric, one that is built around a common OS and that ties their different deployments together into unified security architecture. Fortinet's virtual solutions provide all the same features and functions of our hardware appliances. They run the same network operating system, use the same policy, management, and reporting consoles, receive the same security updates, and apply the same authentication. This commonality gives the customer a lot of flexibility when architecting their network security. This approach to security as a single system, regardless of whether it is implemented as

an appliance or a virtual machine, enables customers to deploy a unified security policy and enforcement layer from IoT to the cloud.

EA: How fast are your customers moving toward virtualization, and do you see a time when they might buy security almost exclusively as software?

KX: The vast majority of our customers adopting virtualization are developing a hybrid deployment approach. While the dynamic workload management that a virtualized environment can provide is increasingly important, so is the heavy lifting of things like Big Data that require high-performance workhorse appliances. For example, for East-West traffic in data centers, or in private or public clouds, the adoption of virtual security is happening quite quickly. The distributed enterprise security requirements are also changing rapidly, especially around the WAN NGFW. Increasing SaaS/IaaS traffic, and the need for localized segmentation have placed higher demands on CPE offerings. Longer term, this may require a virtual CPE offering that cannot yet be virtualized.

EA: How does a world-class security vendor decide which features and functions to include in its platform? Do you do this scientifically, or do you have to make bets on the future?

KX: Predicting the future for security requirements is both art and science. Our vast network of FortGuard threat researchers is located around the globe monitors and deeply analyzes the threat landscape every day. This provides us with threat and trend intelligence, along with awareness of threat patterns and behaviors that enable us to plan effectively for the future. But we also take input from our customers very seriously. They are on the front lines of use cases and requirements, and often request new features through their local sales engineering teams that we hadn't considered. This two-way conversation helps us build products and services that meet the changing demands our customers are experiencing. And finally, we invest in strategic technologies, such as cloud or ASIC technologies, which have a longer development cycle. Happily, the market has shown that these forward-looking investments were the right approach, which has given us a huge lead over our competitors in meeting today's security and networking demands.

EA: Your company has always been focused on securing the enterprise network. As the concept of enterprise network evolves to include mobile, cloud, and virtual infrastructure, is it easy to evolve the focus of your products?

KX: For years we have been developing and evolving our solutions suite toward what we call the Fortinet Security Fabric. This fabric is designed to dynamically adapt to the evolving IT Infrastructure in order to cover its rapidly changing attack surface. It intelligently and transparently segments the customer's network, from IoT to the Cloud, to provide advanced protection against sophisticated attacks. Each security element in the fabric is also aware of each other, allowing them to share

policy and application flow information. This collaborative approach to threat intelligence provides a much faster time to detect, no matter what part of the network is being compromised, as well as the ability to provide a coordinated response. In addition, Fortinet encourages our technology partners to be an integral part of this distributed security framework.

EA: Do you see many differences between the types of threats that your customers are trying to deal with today versus when you first started in this business so many years ago?

KX: Yes. The biggest observation is that today's threats are much more complex, utilizing multi-vector attacks and sophisticated evasion and persistence techniques. Fifteen years ago, the motivation for hackers was mainly notoriety, but today's cyber threats are a multi-billion dollar business, with the intent to extract intellectual property, identities, or monies, or increasingly, they are politically motivated cyber attacks. In addition, today's cybercriminal has access to a marketplace of commercial resources to help them build out their threat lifecycle. The commercial underground provides tools, processes, and even help desk services to assist hackers in defeating security defense systems. It's an arms race, and while a security strategy needs to identify and fend off every threat, the criminal community only needs to get through defenses once. To better protect our customers, the security industry needs to cooperate more in sharing live threat intelligence. Fortinet is a founding member of the Cyber Threat Alliance, which promotes this goal. We also support third party test houses, such as NSS labs and ICSA, to make sure the security systems being deployed provide the best possible protection.

Advanced Support for Digital Investigations

Providing Automated Forensic Tools for Investigators Supports Effective Endpoint Visibility

Patrick Dennis, CEO of Guidance Software

Few techniques in cyber security are as mature, but also as poorly understood as the process of digital forensics. Generally viewed by many industry observers as a "black art," forensic investigation involves experts focusing on providing full visibility and accurate intelligence in computing environments where data has been corrupted, lost, or mishandled. It is a difficult task and requires considerable experience and expertise on the part of the forensic investigator. And yet, with the provision of world-class tools, digital forensic experts continues to exceed expectations, supporting a range of difficult needs including support for tough litigation cases, proactive threat intelligence gathering, and compliance program management.

EA: Has the basic role of the digital forensic investigator changed much in the past decade?

PD: I guess the high-level goals associated with digital forensics, eDiscovery, and endpoint security have not changed dramatically in the past decade, but the specific investigative techniques, the underlying platform used to support investigations, and the consequences of most forensic cases we've seen, have all changed pretty significantly. In our support for the investigator at Guidance Software, we try hard to maintain focus on the fundamentals – namely, providing excellent visibility into target systems and endpoints, supporting accelerated needs to surface useful facts in a super-timely manner, and adding automation to the protection of endpoints. These goals have not changed as we've supported so many companies – almost 80% of the Fortune 100 at last count – in these important investigative and response tasks.

EA: To what degree is the investigator dependent on technology? Is the skill and instinct of the investigator a factor?

PD: That's an interesting question. First of all, every investigator will point to instinct as being critical to a successful engagement, whether it be forensic investigation of a cyber intrusion or an eDiscovery task for some legal matter. In every case, the skills, experience, and intuition of the investigator will play important roles. That said, it is our mission at Guidance Software for the automation to enable data visibility to help discover malware – and this obviously requires advanced algorithms and automation, and to support the entire incident response task from start to finish, managed by workflow process automation. So while everyone agrees that human skill is needed, we believe our platform provides incredible value to the end-to-end process.

EA: Do you see more cyber security teams trying to do proactive forensic investigations, perhaps to find early indicators rather than just waiting to investigate after an attack?

PD: That certainly makes sense, because proactive forensics, which is sometimes referred to as cyber hunting, is really no different if the investigator is searching for early indicators than if the investigator is searching for evidence of a persistent attack. This also helps explain our more recent focus on Guidance Software's tools as endpoint security solutions, rather than as purely investigative support. We realized that since our tools were so good at providing visibility into operating systems, devices, and infrastructure – as the many generations of EnCase users have long know – well, these are precisely the types of requirements associated with good endpoint security.

EA: Has modern litigation changed the nature of digital forensics? I would assume that most lawyers have gotten savvier in recent years about the power of the forensic expert.

PD: Yes, there is no question about it – the legal profession has come to recognize eDiscovery and also digital investigative forensics as essential to their work. We tend to see several primary business areas growing amongst our customer base. First, as you suggest, there is the litigation support that our world-class eDiscovery tools have long assisted. But in addition, we also see significant growth in the regulatory compliance area, breach detection and response in the enterprise, internal employee investigations of insider activity, and as you'd expect – law enforcement investigations, which has always been such a proud part of our business at Guidance Software.

EA: You mentioned regulatory growth. Does compliance have any impact on forensic planning? To what degree, for example, do frameworks like PCI DSS and HIPAA affect or influence the digital investigative process?

PD: They have a dramatic influence on forensics, simply because breaches today involve loss of sensitive personal and business data. These are real business problems and the compliance and regulatory community is determined to reduce the risk. So it should come as no surprise that our EnCase platform would be such an important solution to dealing with compliance problems in the payment card industry, the health field, and many other sectors. We are so proud of our massive reach, with many tens of thousands of trained users of our platform, located around the world.

Advanced Threat Protection for Web-Based Applications

Combining Web Application Security, Real-Time Threat Intelligence, and DDOS Protection into an Integrated Defense

Anthony Bettencourt, CEO of Imperva

Ask any security expert about the top cyber attacks on modern enterprise Web and mobile applications, and they will invariably list technical breach methods such as SQL injection and cross-site scripting, and DDoS attacks on the applications themselves. As a result, it stands to reason that solutions such as Web Application Firewalls (WAFs) will grow in popularity to reduce the risk these common attacks pose. Given the role of WAFs in the application architecture, it also stands to reason that the best cyber security vendors will expand the scope, intensity and applicability of their web-based defenses to protect a customer's entire enterprise and even ecosystem.

EA: Why do you think the WAF has become such a popular device in the typical enterprise?

AB: Web attacks have long been the staple of cybercriminal organizations – according to the most recent Verizon DBIR, 40% of data breaches occur through Web app attacks. A well-designed Web application security strategy, including a modern WAF, is a critical element of an enterprise's overall security posture. In addition, we've all experienced how pervasive Web services and applications have become – think banking applications on smart-phones for the individual and cloud-based file transfer solutions for business. It thus stands to reason that if basic functions like file transfer are done using Web services, then the corresponding security solutions must be designed to protect the shifting and growing portfolio of application services that use these services.

EA: The PCI-DSS standard recommends either WAF functionality or source code review as means for strengthening Web applications. Do you see these techniques as equivalent?

AB: A strong Web and mobile application security strategy will incorporate both WAFs and a review of source code. Our team at Imperva recognizes that these solutions are compatible. WAFs provide a powerful, flexible mechanism for enterprise security teams to deploy real-time protection for their valuable Web services *without having to change the services themselves*. This is an important advance that reduces risk considerably. Source code reviews to improve these same Web services are also a good idea, although unfortunately, the reality is that humans will be unable to catch every potential error introduced into code by another human. We see a large number of attacks on well-documented and old vulnerabilities that just have not been fixed. We have also seen cybercriminals use application-based DDOS attacks that exploit bad application design, rather than just simply trying to flood the network pipe. These kill the enterprise connection and cannot be stopped with network DDOS mitigation tools. Rather, this requires a modern WAF to inspect the application traffic and block these attacks before they can exploit the application.

EA: WAFs began as a hardware appliance inserted to some on-premises segment for inspection of application traffic. How does public cloud usage change this?

AB: You are correct. The modern enterprise is virtualizing, and many are starting to use public cloud services, or considering doing so in the near term. Companies are using the cloud for handling regulated personal information like health records, as well as for storing intellectual property and delivering essential services. This is what cybercriminals target with either DDOS attacks, direct break-ins or through insiders. As enterprises virtualize and consume cloud services, the idea that a solution can reside as a hardware appliance deployed at the perimeter is no longer viable in many cases. Instead, WAF functionality, Web security and nearly every aspect of enterprise security must have the ability to virtualize across all entry and exit points of the enterprise. Today this implies that Web application security is deployed as a cloud service itself, as well as via virtual or physical appliances that can be deployed both on infrastructure-as-a-service clouds and within on-premises data centers. Our platform solution at Imperva, therefore, includes three major functions. First, our Imperva SecureSphere and Imperva Incapsula platforms halt DDOS attacks aimed at Web services, our Incapsula, SecureSphere and Imperva ThreatRadar solutions detect and mitigate live attacks and break-in attempts to Web applications before they can produce damage, and our SecureSphere, Imperva Skyfence and Imperva CounterBreach solutions help discover sensitive data across the enterprise and in both sanction and unsanctioned cloud apps. This includes breach detection of insider access to business-critical assets.

EA: Everyone knows that DDOS attacks at Web services are moving up the stack toward Layer 7 application functionality. Does the WAF offer protection here or do enterprise networks need to complementary DNS and geographic protections of a CDN or similar network architecture solution?

AB: Yes, the layer seven DDOS attacks are more clever and tougher to filter than earlier layer three DDOS attacks that flood a site with traffic. There is no question that defenders have to get smarter, which is why our platform has been designed to deal with live attempts to manipulate Web applications such as those that would be found in a more advanced DDOS attack. We're seeing DDOS attackers use both application and network-level DDOS attacks together. We believe a comprehensive DDOS mitigation strategy needs to be able to mitigate both volumetric and "low-and-slow" attacks against both the network infrastructure and Web applications. Our solutions are designed to deal with these broader spectrum attack scenarios. We combine WAFs with a cloud DDOS mitigation infrastructure that is best able to protect against today's more sophisticated attacks and keep the impact of mitigation out of the enterprise environment, which often are still pretty conventional with Web applications hosted in the data center and accessible through a perimeter gateway.

EA: Do you see more critical functions, perhaps supporting life-safety or critical infrastructure services, gravitating to Web-based solutions? Does this change the nature of the Web application security business?

AB: There is no question that Web services have become the most critical underlying technology for so many aspects of life-safety and critical infrastructure services. The growth in excitement and utility of the Internet of Things (IoT) is a perfect example. While many of the devices in factories or industrial plants today are based on older protocols and access methods, many are being redesigned or extended with Web services. As a result, platforms such as those from Imperva will not only protect traditional enterprise applications from cyber-attacks, but could increasingly play a role protecting essential services. This doesn't change the nature of our business, though, because we have long understood how important our mission is to business, government and the entire Internet ecosystem.

Advanced Cyber Security Capture and Analysis in Real-Time

Security Analytics for Real-Time Security Support Requires Capture and Cyber Processing at Line Speed

General Keith Alexander (ret.), CEO of IronNet Cybersecurity

Protecting the enterprise has long required the use of real-time intrusion detection solutions. The main problems, however, have involved the lack of comprehensive capabilities addressing a range of use cases, products falling short as network capacities have grown, and products not including sufficient advanced heuristic analytic techniques based on sound underlying models able to function, update, and share in real time. By improving in these areas – creating comprehensive capabilities driven by a range of real-world use cases, driving capture and analytics to line speed, and employing embedded analytics based on creative correlative methods – the enterprise can experience much higher rates of advanced threat detection, even if the actor is a nation-state or other highly capable player.

EA: What do you see as the primary gap in cyber security today?

KA: In talking to CISOs around the country, the most common challenge I hear is how to make sense of the dozens and dozens of tools, individual capabilities, and data available to them. What I've heard them say is that they are looking for a comprehensive capability that can take the data they already have and build upon this data to address the tough use cases they face every day, making their analysts smarter and more capable of dealing with threats as they come at them.

EA: How important is it for advanced threat detection methods to operate at high capacity line speed?

KA: I believe it is absolutely vital. Recent reports show that intrusions and operations are moving at increasingly faster speeds, and if we want to mitigate future events, we need the speed to accomplish this as early as possible. That means

not just capturing at line speed, but employing Big Data analytics, also in real-time, at scale, and delivering both analytics results and the ability to mitigate at high speeds.

EA: Do you think it is ever going to be possible to outsmart advanced nation-state adversaries?

KA: I do think it will be possible to outsmart them in tactical or small operations, but strategically this will be very difficult for industry. Indeed, it is fairly unrealistic for us as a nation to expect individual private companies to defend themselves against committed nation-state actors with huge budgets, advanced capabilities, and hundreds, if not thousands, of operators to throw at the problem. As a consequence, it is critical that private sector actors collaborate better with one another, in real-time, and on a 24x7 basis, not to mention that government and industry must work together, again in real-time, as needed in a time of national crisis.

EA: Has it been your experience that the offense, both in government and commercial settings, is getting better?

KA: Yes, the offense continues to get better. In fact, what is most troubling today is the recent increase we've seen in destructive cyber attacks conducted against private industry. In 2012, we saw destructive attacks against Saudi Aramco and Qatari RasGas, including 30,000 computers bricked at Saudi Aramco alone. And more recently, we've seen the destructive attacks here in the United States against Las Vegas Sands Corporation by Iran and against the Sony Corporation by North Korea. Look, the reality is that the offense is not only getting better, it is virtually likely to succeed at some level nearly every time. We need to shift from a posture of not only trying to stop attackers as they enter, but also in finding them quickly once they are in, before they can do real, extensive damage.

EA: What is the role of information sharing in cyber security today? Do countries such as the United States have sufficient legal cover for this to proceed properly?

KA: Information sharing – really threat intelligence sharing – is critical. Private sector companies need to work together, and sometimes even with the government, to truly defend themselves against the range and scope of threats in play today. And the starting point for that effort is the real-time sharing of actionable threat intelligence. While the recently enacted cyber legislation is a great starting point, providing clear legal authority to obtain and share cyber threat information and important liability protection, I think we could do more to incentivize sharing and to ensure that the right people have what they need to defend the nation. It is also my hope that industry and the government will continue to partner together on this effort.

EA: You've probably seen more sizes and shapes of cyber attacks than anyone in the world. What keeps you up at night today?

KA: I worry that our commercial sector and our government are not prepared for large cyber attacks. We fundamentally have not had the key national conversations we need to have about who is supposed to protect the nation from seriously capable, committed adversaries, and how we are going to make that happen. Today, we expect companies – large and small – to protect themselves against all manner of cyber threats. Of course, this is totally unrealistic and unfair to industry – we don't expect them to defend themselves against nation-states in any other context, so why should we expect it in cyberspace? It's not logical. At the same time, while the government is sorting all that out, I think industry can best help itself by right-sizing its cyber security investments and looking at tools and capabilities that can adapt to threat as it morphs and changes.

EA: If the President were asking you today for your best advice about cyber security, and I know this would hardly be a new experience for you, what would you say?

KA: We need to prepare now, and get the authorities to all the right parties. We need to be clear about what we expect from industry and what their authority is to protect themselves. And perhaps most importantly, we need to bring key players across government and industry together and rehearse how we will act in a crisis, including taking action in real-time to protect the nation. And this effort has got to be based on a coherent, intelligent strategy that doesn't set unrealistic expectations or impose government mandates on the private sector.

Embedded Security in Networking Infrastructure

Network Elements Provide a Convenient Platform for Integrated Cyber Security Protections

Sherry Ryan, CISO of Juniper Networks

One of the most familiar tenets of cyber security is that retrofitting security into a design is rarely a good idea. And this is certainly true for enterprise networks. The idea that an IT team might design and deploy a network and *then* design and deploy a security system has come under great criticism of late. The preferable approach involves native integration of cyber security in the network design from the beginning. One of the components of such an approach involves native security protections being integrated into the network elements on which enterprise local and wide area infrastructure is built.

EA: Is the network the most important component in any enterprise security architecture?

SR: The network is obviously the pathway to everything else that comprises an enterprise information system; therefore, it is a critical area for emphasis relative to security architecture. Saying it is most important, when contrasted to the other layers such as infrastructure, application, or data, would be arguable, as effective security is about context between components of an overall architecture. The network – not to mention the overall security architecture – is important, for certain, but the extent of how network security in an architectural context is applied is wholly dependent on a multitude of factors for a respective enterprise, including but not limited to security requirements, risks, data and how applications and information are used and consumed.

EA: What are the security risks of older, legacy networking products in an enterprise environment?

SR: There are unfortunately many unpatched, poorly maintained, and end of life/end of service (EOL/EOS) devices that organizations are unwilling to replace, often from fear that they will disrupt or impact services and productivity. Many legacy networking products therefore do not provide a cohesive mechanism to apply and manage a contextual policy across disparate network components. Instead, they have to be managed to a great extent individually, and one bad firewall rule, misconfiguration or exploited vulnerability can lead to a broad and cascading compromise. Many legacy networking products do not emphasize mitigating risks to the device or to the capabilities it provides. They are oriented instead toward performance and availability (rightfully so), but are missing the mark on protections of the network component itself. As an example, legacy routers have integrated data, control and management planes embedded in the firmware; there is no means to separate them and apply more separation and granular control.

EA: When an enterprise IT team buys a router for use in the enterprise network, is it ever sufficient to just use the native filtering and logging capability of that router for security? Or is it almost always a necessity to build a perimeter network around that router?

SR: It really depends on the enterprise, the risks, threats, regulatory compliance concerns, and objectives. For many organizations that are focused on protecting intellectual property, healthcare or other information that falls under regulatory scrutiny or provide critical services, such as an ISP, ensuring that network devices such as routers, switches, VPN gateways, and firewalls are sufficiently protected is an imperative. Generally speaking, core networking components should be protected and decouple the data, control and management planes via Software-defined networking (SDN), which facilitates much greater granularity of control, and as a result, security, and adaptability.

EA: What is the current trend in threat intelligence management and information sharing for network elements such as routers?

SR: Threat intelligence and information sharing are still, to a great extent siloed. That is, there are many threat information sources and efforts to do better in terms of sharing information, such as TAXII and CybOX frameworks for sharing information between different organizations. But in some ways this information is not necessarily making its way to network elements in a timely manner and often are applied through endpoint or other security specific capabilities like Web URL content filters at the network edge. One issue is the lack of a common “security messaging bus” between network components. Actually, they do exist, but they are proprietary or have not been widely accepted or adopted across network vendors for interoperability. In time, the goal would be to inject IP threat and reputational

information more directly into network elements which in turn are able to share this information across a common and protected management plane.

EA: Do you see network function virtualization as a security challenge or opportunity for enterprise customers?

SR: I see both. It is a security challenge certainly, as it fundamentally changes the way we must think about, plan and deploy network and security architectures, with less focus on physical aspects of applying security and more focus on logical and abstract context across layers. But it is also a big opportunity to rethink the way we have traditionally approached network security, and recognize the potential to more wholly integrate security within the network where the control, visibility and dynamic adaptability are intrinsic to the network itself.

EA: How quickly do you see software defined networking being adopted in data centers and ISPs? Do you see the associated security for SDN as a native capability or will data center and network managers have to overlay an SDN security architecture?

SR: I believe we are on the cusp of a major approach revamp in terms of how we design, deploy, provision, manage, and secure our networks. Virtualization and cloud-based services have profoundly disrupted and directly influenced the need for significant changes to how we approach networking and security. The days of a clear and defendable perimeter are long gone. Enterprises must reorient themselves to this reality, which means that SDN will very likely have to be adopted sooner rather than later or enterprises will continue to struggle with trying to leverage legacy physically constrained architectures that cannot practically support the logical convergence and challenges this poses relative to security, management and control of our data, access, and policies. My hope is that SDN will wholly include security so it can be applied contextually and intrinsically across all network elements. Security is already complex enough without adding another overlay that may not provide the cohesiveness and operational efficiencies needed to adapt to this convergence while ensuring security protections are part and parcel to the network itself.

Modern Enterprise Voice Security Controls

Offering Advanced Voice Security Protections for Executives and Users Desiring Enhanced Privacy

Nigel Jones, CEO of KoolSpan

Most of the attention in cyber security over the past few years has focused on data security controls. And this is understandable, given the number of highly visible enterprise break-ins and data exfiltration cases that have occurred. But the requirement remains that voice communications must be properly protected, and while carriers have done an admirable job improving controls through improved standards, considerable privacy gaps remain, especially for traveling executives. These privacy gaps are best addressed through a combination of encryption, key management, and related security controls for traditional and over-the-top voice security communications.

EA: What are the typical requirements you see from business executives for voice security?

NJ: The requirements we see fall typically into the following categories: Security, user experience, and for some, enterprise features. When it comes to security, people want to know that their calls and texts are protected end-to-end with proven, strong encryption. For user experience, a high quality, easy to use, convenient solution is important. The drawback of most secure communications solutions is that they sorely lack a quality user experience. Our philosophy regarding TrustCall is that if a secure call is as high quality and as easy to use as a regular phone call, then why would anyone ever opt to make an insecure call? When it comes to enterprise features, people ask for a solution that fits into their existing environments, so that it can be easily integrated via APIs into their ERP, CRM, provisioning, MDM, and other systems.

EA: Do you see international travel as a major driver in the voice security marketplace?

NJ: Absolutely. And it almost does not matter in what industry they operate, from finance services to construction, energy, manufacturing, retail, and many others. All international travelers inevitably are targeted by regional actors, whether the local government, business competitors, organized cyber criminals, or even hacktivists. Every international business traveler should assume that everything he or she says in their phone calls, and everything they text to others, will be intercepted and potentially used against them. I can tell you many stories. For example, we have a client whose business development people were talking on their cellphones in a Latin American country about the important bid they were going to submit the next day for a regional contract. It turns out that they lost to a competitor whom, they believe, listened to their conversations, and then slightly underbid them to win the business.

EA: What are the advantages of software-based encryption over hardware? And I guess I should ask the reverse question as well, since hardware has always played an important role in cryptography?

NJ: Historically there was a big difference between hardware and software-based encryption, and the encryption purists argued that a hardware anchor was critical. Today, the reality is that they are converging, in that sophisticated software encryption can rely on other anchors, including the devices themselves and the secure elements of the chips in the devices. At KoolSpan we offer both solutions, and they are interoperable.

EA: Do you see more compliance auditors starting to require voice security in their security requirement frameworks?

NJ: Yes, and it is happening with astonishing speed. Only a few years ago, voice security was a niche market, serving principally government and defense organizations. But two things have expanded the market. First, the cost and level of sophistication required to intercept mobile communications has plummeted. Today, a non-techie can intercept phone calls and texts with equipment that costs less than two thousand dollars. As this intercept cost came down, the volume of attacks has increased dramatically. And second, the global enterprise market is much more aware today of attacks on mobile communications via, it seems, a regular drumbeat of high profile attacks and increasing media coverage. Today, it is fair to say that encrypting mobile communications is a well-recognized best practice and I believe that in the relatively near future it will be mandated by enterprise security teams in all government and business sectors and for organizations of every size and shape.

EA: Have we reached the point where "voice" is essentially synonymous with "mobile?" Or do you still see businesses requiring security for landline voice communications?

NJ: There is no doubt that voice and mobile are becoming synonymous. That said, we do not see landline voice communications going away. For that reason, we offer TrustBridge, so one can make a secure call from mobile into the corporate environment and vice-versa.

EA: What do you see as the role of OTT communications application in the modern enterprise? Will they become more important and will they require encryption?

NJ: It depends on how you define OTT and the various parties involved. We believe that communications will be delivered differently to varying segments of the market. Many TrustCall customers today prefer to implement their solution "as a service," and for them, we provide the TrustCall Global Service, so there is no customer infrastructure or capital expenditures. By the way, we also have carrier partners globally that sell TrustCall to their customers as a service. Many other customers, including some enterprises and most defense, law enforcement, and other government organizations prefer to control their own communications system, so they can protect not only their data, but also control the metadata. This second set of customers will purchase TrustCall DIRECT. We help these customers deploy the necessary infrastructure on their premises or in their private cloud, and we provide training and support, so the customer can manage their communications system directly.

Intelligence-Based Detection of Cyber Threats in Log and Machine Data

Providing an Intelligence and Data Analytics Platform to Detect, Prioritize, and Neutralize Cyber Security Threats

Mike Reagan, CMO of LogRhythm

Just about every enterprise security team has some sort of data collection facility for detecting security events. This can range from a modest log management function in smaller companies to more extensive SIEM deployments in larger companies. But in all cases, optimizing the data management and analysis task to detect cyber attacks quickly and respond to them before they can produce consequences is a difficult defensive task for an enterprise. It becomes even more difficult when the desire emerges to use real-time threat intelligence as a basis for making security decisions.

EA: Do you think that many of the prominent cyber attacks over the past few years might have been prevented through more proactive log management and analysis methods?

MR: It's easy to play Monday morning quarterback and claim that "if they had only been using our technology, they *never* would have been breached." But the reality is that preventing major breaches requires a combination of people, process, and technology. Given what we all have now learned about these prominent breaches, there's no doubt that effective use of security intelligence and machine analytics could have provided earlier visibility to the indicators that threat actors were afoot in the enterprise, and this probably would helped to avoid a material breach or service disruption.

EA: What are the ways in which intelligence, analytics, and SIEM functions can come together in a common platform and solution?

MR: We continuously hear from our customers that they place a high value on having a truly integrated platform. It starts with having all the data, so

comprehensive log management is important, as well as adding endpoint, network and user activity data to the mix. But collecting it isn't enough. The data needs to be normalized and processed so that advanced machine analytics can be applied to it, and so analysts and investigators can perform highly efficient searches and investigations against the data. Once that's in place, an intuitive and optimized user interface is required to surface the intelligence to the right people at the right time so they can respond quickly enough to neutralize threats before they can have a material impact. The latter requires integrated incident response orchestration. Essentially, a full end-to-end threat lifecycle management solution comprises all of these components, and that's precisely what LogRhythm is delivering to our customers through our Security Intelligence and Analytics Platform.

EA: Much has been made of threat intelligence as a major component in cyber defense recently. Is there sufficient access to good intelligence for enterprise security teams?

MR: There are hundreds of threat intelligence feeds that are available, ranging from general commercial and open source feeds, to industry-specific data sources. In fact, there's an overwhelming amount of threat information being disseminated, and much of it becomes outdated very quickly. To make the most of the myriad threat intelligence sources, organizations require a central security intelligence platform that can automate the consumption, corroboration and evaluation of this information with internally generated threat intelligence.

EA: Many enterprise teams have their SIEM in place and have built processes around that tool. Is this sufficient in many contexts, or do CISOs really have to augment the SIEM with more advanced capabilities?

MR: There are thousands of SIEM deployments in place around the globe, but most are first generation SIEMs that were originally designed to collect and store log data or to cull actionable events from basic security devices such as IDS systems and firewalls. Most users of these legacy platforms are overwhelmed by the sheer volume of events they need to evaluate, and are burdened by the complexity and cost of managing the underlying platform. They are also acknowledging that they are blind to many of today's threats because they lack the automated, machine analytics to evaluate the millions or even billions of logs being generated every day. They also realize that simply relying on manual hunting for threats to keep their enterprises secure is not the answer. To achieve comprehensive threat lifecycle management, CISOs are deploying unified platforms that combine SIEM with log management, network and endpoint forensics, and advanced security analytics.

EA: What sort of trends do you see in the cyber security industry today? Are the hackers just growing at a faster rate than the defenders?

MR: We're seeing a rapidly expanding cyber-crime supply chain that's acting as a force multiplier for online crime and cyber terrorism. This supply chain is fueling

innovation at a pace well ahead of the development of technologies designed to keep the bad guys out. In light of this reality, organizations are accepting the fact that cyber adversaries will breach their defenses, if they haven't already. However, forward-leaning CISOs realize that a breach of perimeter defenses doesn't have to result in a material data breach or service disruption. They are evaluating their own organization's security intelligence maturity and focusing on continuous reduction of their mean-time-to-detect (MTTD) and mean-time-to-respond (MTTR) to cyber threats by employing and honing comprehensive end-to-end threat lifecycle management.

Advanced Endpoint Security for Mobile

Extending Threat Intelligence Based Security Solutions for Mobile Endpoints and Apps

Kevin Mahaffey, Co-Founder and CTO of Lookout

Endpoint security – in the form of PC anti-virus programs – was the original computer security solution. And for many years, it dominated the enterprise security landscape with blacklist-based signatures warding off viruses on Microsoft operating systems. Fast-forward to today, and the endpoint has become a mobile device or tablet used to access cloud-based assets that are provider-hosted outside the previously safe enclave of the perimeter LAN. Cyber security protection thus emerges as a challenging goal for these mobile devices and the ubiquitous apps that differentiate each user's mobile experience.

EA: Is mobile security a direct extrapolation of PC security? Or are there truly unique aspects of the mobile endpoint experience that are different from PC usage?

KM: It's funny you should ask that question, because Lookout just announced a partnership with Microsoft to integrate our mobile endpoint security with their enterprise mobility suite. It's a great example of two companies combining their respective core competencies – ours in dealing with the growing threat of sensitive data loss through mobile device threats, and Microsoft's in dealing with the day-to-day enterprise IT requirements driving businesses for the past twenty-five years. Now, to answer your question more directly, I'd say that there are significant differences between PC and mobile security, most of which stem from the architectural differences in how each are used in a typical work environment. For many years, employees would store primary copies of their data on a PC, which led to protection approaches that were focused on the native operating system in a PC. In contrast, modern mobility security assumes that mobile devices will inevitably be

hacked, but sensitive data and assets stored in the cloud will still be secure. Second—and this is a big one—on many mobile devices, employees combine personal and work activities, preventing organizations from employing restrictive approaches such as binary whitelisting or URL filtering that have become commonplace on enterprise PCs. There are other differences, but these are the major ones.

EA: Can CISO teams separate protection of the mobile from protection of the cloud-hosted app? Or do you need an end-to-end solution for users reaching out to the cloud with their smart device?

KM: I believe it's only possible to reason about mobile and cloud security together. Why? From the history of PC security, we've learned that OS security controls and endpoint protection software cannot guarantee 100% of threats will be blocked. Instead, I advocate for a data and application-centric approach, where you enforce conditional access to your cloud data and applications based on risk information received from a mobile endpoint security agent. Of course, you want your mobile endpoint security agent to catch as much as possible on the device, but you cannot blindly assume it will stop all threats. I believe that all organizations will eventually move to an enforcement model where access to cloud data and applications depends on the security posture of the device accessing them.

EA: To what degree will mobile security rely on accurate threat intelligence? And does this require live threat feeds to the mobile device?

KM: I believe that timely and accurate threat intelligence is the most critical aspect of any endpoint cyber security solution. In fact, you could probably say that good intelligence is the most critical aspect of *any* cyber security solution. However, effective use of threat intelligence isn't simply getting feeds of tactical threat indicators such as application binary hashes, domains, and IP ranges and using them to trigger alerts. The reason for this can be seen in the challenges that emerged for the signature-based blacklists used for many years on PCs. I don't have to summarize here all the issues with signatures, but suffice it to say – everyone soon realized that a different approach was needed. Effective threat intelligence includes two parts. First, a strategic, qualitative intelligence that helps an organization prioritize limited resources against their highest risk threats. Second, tactical threat indicators that can be used, not just alone, but also to add context to machine learning and big-data correlation systems to protect from zero-day threats quickly and reliably.

EA: Why do you think the compliance auditors haven't been more aggressive in demanding mobile endpoint controls?

KM: I think that auditors, regulators, and compliance managers absolutely are becoming more aggressive in demanding mobile security protections. Granted, the

majority of major attacks to date have occurred with the usual assortment of PCs, physical servers, enterprise-hosted systems such as Active Directory, and Internet systems such as DNS. However, everyone on the planet knows that cyber attacks are moving in the direction of mobile, if only because that's the direction IT is moving. Skilled hackers know full well that if you want to create problems for an enterprise, go after their mobile devices. Now, translating the mobile threat into compliance standards is not straightforward, because compliance standards are typically developed as a trailing indicator of major threats that have occurred in the past. This helps explain why compliance frameworks are often insufficient at stopping new attacks. So this will have to change, but as I implied earlier, I think we are moving in the right direction today.

EA: Will we ever see a global attack where users have to rush to turn off their mobiles to avoid losing their data or apps?

KM: I hope not, but an auto-propagating worm on mobile is a distinct possibility. Such a worm would likely rely on some form of application or device exploit. On PCs, enterprises can respond to exploits by pushing patches and mitigations to endpoints or blocking malicious traffic on the network. On mobile, however, enterprises do not fully control a device's firmware update process and can only control the update cycles of enterprise-managed apps. Further, enterprises cannot rely on network mitigations because mobile traffic does not typically traverse the enterprise perimeter. On the positive side, mobile device operating systems have stronger built-in security than PCs, making it more expensive to build such a widespread attack. The best advice I can offer is that deployment of mobile security should become a greater priority so that, if such an auto-propagating worm does occur, organizations will have visibility into which devices are vulnerable and the ability to mitigate the threat.

Threat Visibility Across the Enterprise

Offering Real Time Network Situational Awareness Through Enterprise Network Infrastructure and Asset Discovery

Reggie Best, Chief Product and Marketing Officer of Lumeta

The cyber security community has certainly seen its share recently of high tech solutions to the cyber security threat problem. Machine learning, behavioral analytics, and artificial intelligence are example techniques being applied to the enterprise in the hopes that threats can be identified. And yet, if you ask any practical minded security expert what the real issues are on their network, they will generally point to common sense issues as the bigger problem – with accurate inventory as perhaps the most commonly cited concern. Knowing what is running on a network provides the type of situational awareness base that is essential to supporting any type of enterprise security program.

EA: Do most IT or security teams have a good understanding of the things connected to and running on their network?

RB: I am alarmed by how little authoritative understanding many IT network and security teams have about the current state of their networks. This is certainly true of large, physical, static networks, where one can still point to pretty much all of the actual infrastructure sitting in datacenters or closets. In those environments it's customary for 20% or more of assets identified in a competent scan to be new to the IT team. And that visibility challenge is exacerbated in modern networking environments by even more dynamic change involving mobile endpoints, virtualized compute and virtual network functions, private and public cloud, and the emergence of software defined everything. Those scenarios are even more opaque as infrastructure and assets can be transitory – here today, gone later today. The

best scanning tools will miss all of that activity. Some academic research indicates as much as 40-50% of endpoints are being missed by periodic vulnerability scanning with the proliferation of mobile. Something real time is needed from a visibility perspective.

EA: Does a lack of visibility into an enterprise network create problems for handling live incidents?

RB: Absolutely. Lacking enterprise network visibility is like trying to prosecute a ground war without a precise understanding of the terrain. If you don't know what the infrastructure is, where the edge is, what endpoints you have, how they're connected, and what routes are present, then there's no way to quickly and efficiently handle incidents that may lead to a breach. Mostly, you won't even know the incident happened. Visibility, monitored and provided in real time, needs to be the foundation upon which a modern security program is developed.

EA: Lots of CISO teams talk about situation awareness today. Do you see network discovery, scanning, and inventory as critical tasks to support this objective?

RB: Yes, these activities are table stakes. Fully indexing what you have via a combination of passive monitoring plus active interrogation, such as putting probe packets on the wire to hunt actual pathways, are the best methods of achieving situational awareness.

EA: With the shift in most enterprise networks to cloud-based virtual operations, how does this affect the task of performing network discovery and situational awareness?

RB: The biggest challenge we've seen with cloud and virtual operations is the need for real time. Often these instances may be silos with very dispersed administrative rights and access. Network and security teams have limited control over what may be happening in these shadow IT instances. We once had a very technologically savvy customer who was doing DevOps in the cloud. They used a VPN from their enterprise network into a virtual private cloud (VPC) instance provided by one of the well-known public Infrastructure as a Service (IaaS) cloud providers – a very common configuration and scenario in the enterprise now. At one point, the security team identified a new virtual machine being spun up in the IaaS VPC, which by itself isn't a bad thing, as developers are working with instances up there. However, when actively probed, shortly after it came online, that particular virtual machine was acting like a packet forwarder, decrementing TTL like an IP routing element. Following on with an automated leak path analysis, it was found that the VM was forwarding traffic outside the VPC instance to the Internet. The edge of the enterprise network had been changed dynamically. But this was only visible for the times when that VM snapshot was run. Without real time, the evidence of that activity would have been very difficult to reconstruct. And someone very expert

would have had to be involved in looking for it forensically, after the fact, through logs from various systems.

EA: Do you see a shift in inventory management and discovery from a hardware asset focus to a more software-oriented focus?

RB: While hardware inventory and discovery isn't going away, there is absolutely an inclination towards more assets being virtualized. Those certainly include virtual server instances, endpoints, or hosts, running various operating systems and applications. Increasingly, virtualized and software defined network functions themselves – like routers, firewalls, switches, load-balancers, and proxies – are being utilized within enterprises and services provider environments. Inventory management and discovery needs to keep up by being real time sensitive.

EA: How important are visualizations in demonstrating network inventory, asset management, and situational awareness?

RB: Our customers are certainly keen on the rendering of indexed information and anomalies in a visual way. Especially with large and complex networks it takes too long to identify problems by sifting through tables with hundreds or thousands of line items. We certainly spend a lot of time on maps, charts, graphs and other rendered visualizations that involve beaconing, color changes of nodes or edges, and graphical illustration of outliers to highlight where an anomaly may exist. The objective is faster time to detection, repair and remediation. Visual representation of that information is pretty important.

Reducing the Risk of Malware in Software

Reports of the Demise of Anti-Virus Solutions are Incorrect as Newly Effective Security Techniques Emerge

Marcin Kleczynski, CEO of Malwarebytes

The original computer security solution for the PC involved the use of anti-virus signatures to clean up viruses, Trojans, and worms. The technique worked for years, until the bad guys figured out how to create variants that sidestepped signature patterns. Many CISO teams have hence reduced their focus on antivirus, but the good news is that a renewed emphasis in the community has led to improved methods for handling malware. As a result, the technique should experience resurgence across endpoints including mobiles and virtual machines.

EA: Does traditional antivirus software detect and mitigate malware?

MK: No, that's exactly why I started Malwarebytes. Symantec, McAfee – all of these companies built their engines and core technology in 1985 and it's my view that they simply haven't updated them much since. At the same time, criminals are getting more and more sophisticated and agile. Most traditional anti-virus software simply cannot keep up with today's generation of cyber criminals anymore because they are being outpaced and using signatures.

EA: It seems like the idea of any type of signature has really gotten a bad name in the past few years. Is this a fair characterization?

MK: Signature-based anti-virus software definitely has some flaws. Using new techniques like advanced heuristics, machine learnings, and behavioral monitoring will get you a lot further than signatures.

EA: Does behavioral analytics play an important role in the detecting of malware on a system?

MK: Yes, absolutely. Addressing gaps in security related to user behavior *before* you are actually targeted is crucial. Behavioral rules are a core component of how our technology catches instances of zero-day malware and the exclusive technology of our anti-ransomware.

EA: What trends have you seen recently in malware design? I assume the malware is getting better.

MK: Ransomware will soon be the most used type of malware we have ever seen. For example, we see almost 5 to 10 new ransomware types created every day, each a little different. This will not slow down. In fact, we will see double to triple the numbers that we have seen with any other type of infection in the past.

EA: Is there an appreciable difference in detecting malware on PCs versus Macs – or between PCs and mobiles?

MK: The ability to detect malware on a Mac is mostly just hindered by a false perception of risk and a resulting lack of preparedness. Many Mac users are genuinely puzzled when they learn that they have been infected, as they believed that Macs were "immune." On forums and blogs, people often tell Mac users not to install any kind of anti-virus software. This leaves users in a difficult situation when they get infected and are trying to find a way to solve the problem.

EA: Do you think we'll ever see a time when operating systems and applications can be free of the malware risk entirely?

MK: Probably not. And even if we do, cyber criminals will find a new way to infect our systems. Their attack vectors are always evolving and as they do, we will continue to fight the most dangerous threats out there.

Providing High-End Cyber Security for Cloud Apps

Cyber Security Solutions Need to Adapt to the Modern Enterprise's Shift to the Cloud

Sanjay Beri, CEO of Netskope

One of the major goals of any application provider, whether in the enterprise or consumer marketplace, is to ensure that all required applications are convenient for user access in a secure manner. This goal was easier to meet in the context of private hosting, but with the need for ubiquitous public cloud hosting, security and compliance now require a new approach – best done in the context of a cloud-deployed protection capability, accessible over the Internet, but carrying all the benefits of a virtual private network.

EA: What are some of the cyber security challenges CISO teams face when it comes to cloud apps?

SB: First, even though there's greater awareness today than ever, CISOs are still mostly in the dark about cloud app usage in their organization. Our experience suggests that many still believe there are fewer than fifty or so apps in use, when in fact the average we've measured is 935. This is not surprising when you consider that entire corporate functions are doing business in the cloud, including HR, finance, sales, engineering, and marketing. The second challenge is determining whether there's risk, and this begs several important questions: Are apps being used for exfiltration of intellectual property? Is regulated data being uploaded to unsanctioned or un-secured cloud apps? Answering and addressing these questions is an important concern. The third challenge involves protecting data and ensuring compliance in those apps. Since over half of all cloud traffic originates from outside of a browser and from remote users, it's hard to put contextual policies in place that do the practical and necessary task of preventing exfiltration of sensitive data or governing activities like share, upload, edit, download, delete, and so on. The fourth and final challenge is stopping or responding to threats that propagate in the cloud.

Popular apps like Box, Dropbox, Google Drive, and OneDrive are being used to infect users, spread malware, perform command and control functions, and host exfiltrated data. Because many companies don't inspect the SSL traffic of many popular apps, it's easy for malware to fly under the radar of traditional tools like Web proxies and firewalls. These are all challenges that cloud access security brokers (CASBs) address.

EA: What are the deployment challenges CISO teams typically have in making cloud-hosted applications available to their users?

SB: Many of the features that make cloud-based applications so attractive, such as sync, share, and ease of collaboration, are the very things that put corporations at risk when it comes to cloud usage. But rather than not allow users and lines of business to deploy those applications, which is not an option, information security teams need to enable them. This means putting the proper admin controls in place, complete with separation of privileges, ensuring that only authorized users on appropriate devices have access, and that organizations can govern specific activities that users can and can't do. This also includes things like sharing, uploading, downloading, editing, approving, deleting, and more. These are often decisions that get thought about after the application deployment takes place.

EA: What do you see as the best approaches to steering cloud app traffic to users? Does this require coordination with the ISP?

SB: When it comes to steering traffic from the user through a policy enforcement point and then to the cloud app and back again, flexibility is key. Organizations need to satisfy any number of cloud security use cases now and in the future, which requires a variety of deployment modes, from proxy to API to secure TAP mode to log-based discovery. At Netskope, we call this a multi-mode architecture. Using a combination of modes enables you to get the most out of your cloud security provider or CASB. For example, one powerful technique involves using information garnered in e-discovered sensitive data in a sanctioned app using the app's API to inform real-time, in-line policies to catch data exfiltration to unsanctioned apps in proxy mode. In fact, three-fourths of Netskope customers deploy in multi-mode. In-line deployments require global points of presence and peering relationships with key service providers. This is something you should require of your cloud security provider or CASB.

EA: What are the pros and cons of doing cloud app security as an on-premises device versus a cloud hosted capability?

SB: Several large, regulated organizations – this includes some of the largest banks and energy companies – have chosen to deploy cloud cyber security on-premises or in hybrid mode, with some portions remaining on-premise, on their journey to the cloud. Some view this as an interim step and others view it as longer-term. For

organizations that process user data and are beholden to strict privacy and data residency requirements such as ones serving European Union customers, the on-premise or hybrid model may be the right choice. That said, many organizations have chosen the cloud model because of its inherent benefits and flexibility, and then addressed their privacy and data residency concerns by choosing cloud providers with in-region cloud locations and the required security and privacy certifications. The advantage of deployment flexibility is you can enable any use case, irrespective of whether users are on-premise, remote, or mobile, all without having to hairpin cloud traffic back to the corporate network.

EA: This might seem like a dumb question, but do you find that most CISO teams even know what applications are actually available to their users in the cloud? And yes, I guess the question also extends to whether most CISO teams even know what applications that they have?

SB: They are coming around to the fact that there are a lot of applications in use in their organizations, but don't yet appreciate the sheer number, breadth of usage, amount of data, level of spend, and creation of data silos and complexity. Many information security teams think of cloud in terms of Dropbox and Twitter, but don't realize that HR is using cloud apps to onboard and track employees; Finance is using them to authorize payments, do payroll, and visualize key business metrics; Development is using them to build product, collaborate on roadmaps, and manage bugs; Marketing is using them to generate leads and develop pipeline; Sales is using them to track and close deals; Customer Support is using them to measure customer projects and gauge satisfaction; and on and on. Quite simply, cloud apps are the way we work today.

EA: What sort of telemetry and details are useful for security teams to collect regarding cloud app usage?

SB: In order to truly gauge risk and govern usage in a smart way, information security teams need to collect and correlate security metadata about the who, what, when, where, with whom, and with what content of all cloud transactions whether in a sanctioned app or unsanctioned one. For example, they need to know if anyone in the AD group *insiders* has shared sensitive financial information outside of the company; or if anyone in Customer Support has downloaded customer data from any CRM app to a personal device; or if any unauthorized developer has uploaded a workload to an IaaS; or whether anyone outside of the HR team has edited salary data; and so on. By combining and correlating details about the user, group, attribute, device, OS, device classification, location, app, app category, app risk, content type, content profile, recipient, and recipient's company, security teams can get a deterministic view about whether a security incident really occurred and what the steps are in the audit trail before and after the event. Beyond forensics, they also need to incorporate those same granular details into controls to enforce policies such as "No access to Office 365 for BYOD devices," "No sharing if you are an

'insider' and the recipient is outside of the company," and "No editing financial data if you are not an authorized user in the Finance group."

EA: You are a cyber security industry veteran, so what sort of trends do you see in the protection of enterprise assets? Are things getting more secure?

SB: In the history of information security, threats have always followed users and data. I know that sounds simple and obvious, but it is important. Enterprise data are moving to cloud apps and environments more than ever, and as a result, enterprises now have hundreds of virtual, shadow IT departments. Similarly, users are performing the bulk of their transactions outside of the corporate perimeter, doing their work remotely on mobile devices, going directly to the cloud or Web without transiting their corporate firewall or security infrastructure. There is no reason these natural trends should make enterprises less secure. Enterprises must shift their thinking from ports and protocols to decoding APIs – the language of today's cloud and Web, that articulates user activities like share, download, edit, and more. Similarly, they need to use that understanding to develop security policies that enable apps while curbing risky activities versus blocking apps altogether. What's required is a new breed of security solution that is API-aware, cloud and remote user capable, and can enable the business in a more nuanced way versus forcing IT down the binary path of allow or block.

Ultra-High Performance Network Security Analytics

Maintaining Full Packet Data Capture in the Presence of Significant Growth in Enterprise Network Size and Scope

Dr. Parag Pruthi, CEO of NIKSUN

In the presence of cloud transformation across all enterprise networks, it is easy to forget that massive capacity growth continues across large segments of the global network infrastructure. And while many enterprise networks continue to scatter their perimeter chokepoint, which one might presume would reduce the need for high capacity and performance, the sheer volume of data traversing the global infrastructure continues to drive the need for not only high performance data capture and analytics, but ultra-high performance support for such critical security functions.

EA: Does it surprise you that network capacity and performance needs have continued to grow in the presence of enterprise perimeter redesign?

PP: It doesn't surprise me, because the vision of the designers of the Internet is actually only now being realized. Remember that the original idea of the Internet was to give users with terminals access to remote computing, and while service providers worked hard to handle high capacity needs, endpoints and servers in the local data center soon grew much more powerful. As a result, large-scale data rarely moved beyond the perimeter. As fiber optics unleashed bandwidth at the core, driving carrier backbones up to 100Gbps, redesigns started to occur where one was no longer limited to local computing, but computing could be spread across the Internet and this gave rise to a proliferation of thin mobile smart phones, cloud computing, and real-time services like video. The demand for more bandwidth in the last decade has led to a complete makeover of the structure of the Internet, from a pronounced hierarchy to a flat structure that blurs the line between network edge and core. During this evolution, the need for security has increased as today's

enterprises are faced with more sophisticated and damaging attacks to enterprise data centers, mobility, cloud, and IoT. And with the proliferation of Big Data, data centers handle more traffic, which fuels the need for ultra-high performance support for critical security functions such as loss-less packet capture and support for analytics.

EA: How fast can a packet-capturing engine go in trying to keep up with a massive network load?

PP: When our team here at NIKSUN first started developing solutions for high-speed capture, we met the challenge of economically supporting 100Mbps, even though most observers said it could not be done. We had similar success at 1Gbps and 10Gbps networks, in spite of skepticism that such packet rates could be handled. Our success was driven by a commitment to building a robust solution that would take full advantage of Moore's Law and that would employ techniques such as stream computing, parallel processing, multi-threading, data base management, and Big Data analytics. The result was our Supreme Eagle product, which is a single-unit modular hardware platform, engineered with the latest high-performance processors, ultra-fast memory, and NIKSUN's next-generation core IP integrated within its own interface line cards. It ensures full line rate data capture and processing ranging from 20 Gbps to 100+ Gbps without dropping a single packet and supports storage up to 10 PB. In comparison with existing industry offerings, Supreme Eagle requires considerably less rack space and power consumption, and delivers more processing capacity and storage than comparable solutions.

EA: Do you see corresponding advances in the quality of network security analytics? That is, as more data is captured at higher rates, is the security analysis still any good?

PP: Security analytics can be compared to fishing. That is, if there are many fish per unit volume of water, then one simply casts a net and pulls out many fish. Similarly, if there are many security issues in a network, then it is easy to cast a security net and identify a large number of threats. But if the density of fish – or attacks in a network – is relatively low, then more intelligent methods are required, and that is where network security analytics comes in handy. For example, advanced cyber attackers, like clever fish, know what traps and signatures are set, and can devise mechanisms to evade security. This can include taking advantage of the distributed nature of a typical large area network. So at NIKSUN, we distribute the analytics, scaling it to big and small nodes across multiple passes, not to mention supporting analytics on virtual machines. An additional problem involves keeping up with attacks embedded in large networks in real time. The NIKSUN team approached this problem by not rushing into the market with a point solution, but rather developing real-time management solutions that focus on streaming data, distributed across multiple sensors. For such distributed streaming data, popular frameworks such as Apache Hadoop that require the different streams to be transferred to a single location for centralized storage and analysis of the data are simply no longer an

option. Instead, we focused on supporting efficient analysis of the hyper data our product generates. The basic challenge posed by such distributed streaming data is how to mine and analyze highest-quality traffic data that is collected in different geographically-dispersed locations and is made available to the analysts in two basic forms — as high-velocity sets of streaming data to be used for real-time analytics and as high-volume, static and highly-structured datasets to be used for network forensics and back-in-time analysis.

EA: Do you see any differences in the way network service providers try to capture data as they shift their networks to SDN?

PP: With SDN, you can equip switches with predefined functions in hardware and let the controller select them for different measurement tasks. In this sense, there will be a role for SDN with respect to traffic monitoring, but this role will be likely quite limited because switch hardware remains prime real estate and severely limits the monitoring tasks that can be performed in the data plane. In particular, lossless packet capture at Gbps line rates for network forensics and back-in-time analysis is a monitoring task for which SDN is ill-suited. On the other hand, SDN can play a dominant role in data analysis for real-time control decisions. In fact, when coupled with an SDN controller, network data analytics can detect and identify nefarious traffic in (close-to) real time. For network security, such solutions enable the timely detection of network attacks, followed by swift and timely mitigation. It is in this role where I see SDN playing a critical role as network providers embrace SDN and reap the benefits offered by a programmable data plane in general and programmable switches in particular.

EA: What are some trends that you are predicting in the coming years? For example, do you see enterprise security teams ever getting to the point where they can actually stop advanced attacks from nation-states?

PP: I see increased investment, but I also see considerable hype. This can include reports that autonomic defensive systems are just around the corner, or that deep learning-based artificial intelligence will revolutionize cyber security. The reality is that the holy grail of cyber security – namely, the detection and reverse engineering of attacks that have never been seen – will continue to require human involvement. NIKSUN's Supreme Eagle supports automation as appropriate, but also enables the domain expert to interact with whatever data is needed to learn about the unknown. If the right balance is struck between using self-learning systems and domain expert-driven discovery and exploration, then future enterprise security teams will be able to detect attacks from nation-states and stop them in their tracks.

Implementing End-to-End Cyber Security Solutions

Comprehensive Security Solutions Assist Buyers in Dealing with the Complexities of the Marketplace

Dan Burns, CEO of Optiv

Providing cyber security solutions to the enterprise requires the ability to combine information security solution design skills with deep knowledge and expertise of available cyber security products and solutions. Add to this the requirement that the solution provider understand the client's business domain, as well as its internal and external processes for purchasing, billing, maintenance, upgrade, patching and many other organizations activities. All of these skills and required support capabilities roll up into the modern cyber security solution provider. And with organizations evolving their enterprise networks to cloud-based SaaS usage, the need for trusted solution provider capabilities will only grow.

EA: Is the cyber security industry – with all its vendors, compliance requirements and need to support changes in IT – more complicated now than it ever has been?

DB: I think you could make the case for that position. Security teams can no longer centralize their security functions into a perimeter gateway, like so many businesses did in the mid-Nineties. It was simpler then, with the firewall and its adjacent security products like IDS and simple DLP all bundled into either one product as in a UTM, or connected together in a small cluster, as in a typical DMZ. Add to that some anti-virus on the PCs and you had the vast majority of enterprise security architectures for over a decade. The role of the value added reseller (VAR) in those days was to ensure that the IT security team was getting the best possible deal. And in many cases, the VAR would manage the contracts, purchasing, billing and so on. Today, however, the typical security architecture combines so many different control areas from advanced identity and access management, to security and behavioral analytics, to adaptive authentication, and on and on. With this evolution

in architecture, the traditional VAR has also evolved to a modern security solution provider like our team at Optiv, offering far more value than just helping to arrange and manage deals with vendors. The modern security solution provider is a trusted partner.

EA: When one of your clients is looking to simplify this security complexity, what are some things you typically recommend?

DB: First of all, security teams should make sure to buy the correct solution for the existing threat in their specific environment. Avoiding purchase of products that might be part of a current fad will only complicate matters, and a good security solution provider partner can help sort this out for clients. Furthermore, with the progression to virtualization, it is now possible to implement so many more functions without the need to purchase hardware. This simplifies procurement, provisioning, billing and maintenance.

EA: Let's stay with that topic. Do you see most IT and enterprise security teams trending toward cloud solutions over traditional enterprise hosted hardware? And are they really using public cloud service for their applications?

DB: Some industries are more aggressive than others, but there is a clear trend toward virtualization and cloud – and that includes SaaS applications hosted in multi-tenant public clouds, a decision that might have been totally unheard of just a year or two ago. It's the economics that drive this decision in the data center, and where IT security teams used to fight the trend, citing concerns about data handling and operation security, now they are investigating new solutions such as cloud compliance and cloud access security brokers.

EA: What are some of the global product and service trends in cyber security? Do you see more organizations buying locally to avoid concerns about code or system integrity?

DB: We see just the opposite in the security marketplace. The landscape of products and services for dealing with cyber risk is now a vibrant global assortment of vendors from so many different countries. Take Israel, for example; the wide range of options from new start-ups in that country us staggering, and our clients are looking for ways to extract that value into their enterprise. So while it is reasonable to worry about the lineage or legacy of a product that might have been developed in another country, it is probably not reasonable to carry those concerns to the extreme of only buying from local domestic vendors. Our clients buy solutions globally.

EA: With ISPs and data centers moving to software defined networking with their on-demand capabilities, what will be the role of the solution provider if a client can just point and click to provision new services?

DB: Well, the value brought by solution providers goes way beyond support for provisioning, so we welcome the self-service aspect of many cloud applications and SaaS offerings. Our role will grow considerably in this new virtual environment as our customers require trusted partners to help sort out the different deals that are provided by different vendors, hosting providers and security service providers. Our team at Optiv is so excited to step up to this challenge and we look forward to helping our clients' transition to newer and more secure enterprise architectures.

Redefining Next-Generation Cybersecurity Platforms

Providing Next-Generation Security Protection Requires More Than Bolted-On Features

Davis Hake, Director of Cybersecurity Strategy, Palo Alto Networks

Weaknesses in traditional firewall platforms have now led to the development of next-generation platforms that provide more automated protection. These next-generation solutions include application-aware processing, as well as embedded integration of related features, such as intrusion prevention. But with enterprise attacks via APT continuing to occur on a daily basis, CISO teams need their next-generation platform provider to reach even deeper in term of technology innovation to help integrate advances in threat intelligence, virtualization support, mobility enablement, and public cloud use.

EA: How would you characterize the main differences between a next-generation platform and a more traditional one?

DH: When companies began to leverage the Internet for business purposes back in the 1990s, the prevailing approach to cybersecurity became the concept of "defense in depth." The idea was that network defenders would deploy multiple prevention and detection controls within their networks and hope that one or more of them would prevent successful cyberattacks. Typically, the traditional security platform included, at least: deploy a firewall, an intrusion prevention system and an antivirus system. Some organizations deployed more controls, and some deployed multiple versions of the same control from multiple vendors. In other words, there might be an intrusion prevention system from one vendor and another intrusion prevention system behind it from another. Initially, this approach worked fine but as the cyber adversary has grown more sophisticated, and as networks have become more complex, defense in depth has proven to be less and less effective. Security teams now operate dozens of siloed products from multiple vendors, which constitutes an

onerous management task that can overwhelm limited human capital. At Palo Alto Networks, we are focused on delivering a next-generation security platform. This means having natively integrated and centrally managed technologies that enable consistent security posture from the network, to the cloud, to the endpoint. The platform also automatically creates, disseminates and ingests protection mechanisms against new threats. This puts our customers in a position to prevent successful cyberattacks.

EA: Do you see next-generation firewalls moving in the direction of a more distributed, virtual perimeter?

DH: There is no question that with virtualization, cloud and mobility initiatives, the perimeter must become more distributed. Your data needs to be protected regardless of where it lives. The use of any public, cloud-hosted application, for example, stretches the perimeter out from the gateway to the hosting environment for your virtual app. Similarly, with the distributed cloud workloads that come with any modern application architecture, such as a content distribution network (CDN) serving as the basis for live collaboration applications, a distributed perimeter is the only reasonable solution.

EA: Why do you think – with all the advances we've seen in cyber security in the past decade – more and more successful cyber attacks? Shouldn't we be seeing more success?

DH: Despite ever-increasingly powerful security technologies, fundamentally this is an economic problem. As the cost of computing power decreases and attacker knowledge proliferates, it becomes cheaper and cheaper to launch successful, automated attacks. This provides a huge profit motive for attackers while increasing the cost for defending against them. In response to this, the security industry has largely walked away from trying to stop attacks, toward providing detection and response capabilities. Unfortunately, we are now seeing the effects of this play out in the headlines. In contrast, the core strategy for Palo Alto Networks has been to focus on prevention first. By automating the integration of network defenses, we can reduce the likelihood of an attack's success at multiple points along an attack life cycle. When you then share data on how to inoculate against an attack with a global network of systems, this raises the cost in both time and money for an attacker and, over time, shifts the economics in the favor of the defenders.

EA: What advice do you have for CISO teams who are trying to figure out how to improve, remove, or change their perimeter?

DH: That will really depend on the architecture. Many of our customers have traditional perimeters with all their critical applications hosted inside, usually in private data centers. These companies will need to take small steps toward cloud, perhaps by virtualizing the private data center first. We've virtualized just about

every one of our products at Palo Alto Networks to support such activity. But for customers who are more advanced, the transition is a bit easier. In every case, however, I'd give the advice that they should make sure to work with vendors who understand cloud.

EA: What type of malicious actor worries you the most – is it the nation state, the terrorist, or something else?

DH: It's all of the above. Sure, nation-state actors are going to be better funded and more determined. But we've also seen young people, the ones we used to call script kiddies, doing pretty amazing things. Just look at the presentations at conferences or the success of bug bounty programs. The techniques being demonstrated in such areas as connected cars, IoT and mobility are pretty advanced and have a huge capacity for damage as ransomware has shown us recently. So all types of malicious actors are converging on common capabilities, where you can't just worry about one specific threat anymore.

EA: Are you optimistic about the future of cyber security protection of critical infrastructure? Or are we headed for a Cyber 9/11?

DH: It would be unrealistic to say that a massive attack can't happen, because every security expert in every country knows that it could. But what we have actually seen happening is the slow bleeding of public trust. A recent study by the NTIA claimed "Americans are increasingly concerned about online security and privacy at a time when data breaches, cybersecurity incidents, and controversies over the privacy of online services have become more prominent." To begin to reverse this trend, we all have a role to play in understanding that, while asymmetric threats are tough to predict, just like terrorism, you can do lots of things to reduce their risk. And with new innovations in technology, efforts toward community information sharing and truly improved virtual, distributed architectures emerging around the world, I think we are beginning to see some roadmaps for success over the long term.

Advances in Enterprise Identity and Access Management

Establishing Identity and Access Management as a Primary Control in Premise and Cloud infrastructure

Patrick Harding, CTO of Ping Identity

One way to differentiate cyber security products and services is based on their function as either a defensive protector of attacks (like an IPS) or as an aggressive enabler of new services (like a secure payment service). The challenge with identity and access management is that it must be both; in fact, if there is one cyber security technology that has the most interaction with, and dependency on, enterprise IT infrastructure, it would have to be identity and access management. CISO teams therefore must get this aspect of their program right, and cannot forget how important the function is to both stopping attacks and enabling business.

EA: In your experience, is there any aspect of an enterprise security program more important than identity and access management?

PH: Given our focus and passion at Ping Identity, I guess that's an easy question to start with, because clearly we see the growing value of identity and access management in the enterprise for cyber security. For many years, this was considered a back-office component in IT, with the registration and maintenance functions viewed as including less attractive work such as the mundane, day-to-day administration of user accounts and access to business applications. Even the user support functions were often relegated to bureaucratic teams who would place you on hold for an hour if you needed to reset a password. The cloud, mobile computing, and the rise of APIs have all contributed to making IAM critical for virtually every customer we deal with.

EA: What role does the cloud play as either an enabler or detractor from the identity and access management goals of an organization?

PH: It is certainly the cloud that has made identity and access management the new primary control in enterprise security with cloud and SaaS applications now being used on par with private cloud and on-premise applications. The Ping Identity Platform gives users quick, secure access through a common set of federated identity and access functions that provide multi-factor authentication, single sign-on, access security, intelligence and analytics, directory, and provisioning. All these functions have to look like a seamless interface, but also must provide a uniform level of security control for cyber security teams and auditors.

EA: So many teams have had colossal failure trying to extend their identity and access management program to larger contexts. For example, failed programs to merge identity and access often follow corporate mergers. Similar failures have been seen trying to do this with cloud. Why so much difficulty getting this aspect of a security program right?

PH: It's all about managing the complexity of these projects, and at Ping, we've tried to create a platform with simple, easy-to-use features so that federation, cloud integration, and other common failure points can be easily managed. Take multi-factor authentication, for example; it is no secret that employees are pretty tired of the non-uniform management of passwords for accounts, systems, networks, and other points of access. This is further complicated by the use of one-time passcodes, hard token, and biometrics. We have simplified this through an MFA service that uses a mobile app authenticating users with a swipe or touch on their self-registered device. Although, we can't change decisions organizations make about authentication policy, we can provide a flexible platform that can make it easier to support those policies. Single sign-on is another good example of a solution that improves productivity for users, improves security, and simplifies complexity in this area, especially with the progression of enterprise identities to cloud applications.

EA: In your estimation, is the identity and access management function best positioned with the CISO or with the CIO?

PH: We see both, and the truth is that the local staff in each area will play important roles. Maybe it's less important which organization managed the identity and access functions, and more important that the correct set of individuals with the right backgrounds, funding, and support should have this responsibility in the enterprise. By the way, as a new primary control, it is also true that identity and access management become a more distributed responsibility across the entire organization. Internal and external audits, for example, almost always have identity and access as either a finding or a recommended control improvement, so no enterprise group can just hand off these important functions to a single group and have them take care of matters. Everyone in the organization must participate in making the identity and access infrastructure and set of services work in a way that is secure and enabling of the local business requirements.

EA: Do you see federation models continuing to grow? For larger organizations, the federation model is beginning to look like a rat's nest of distributed trust. What do you see in the future?

PH: It will grow as long as the supporting platforms for identity and access management are maintained in a simple, well-designed manner. Our support at Ping for Microsoft Office 365 and Google Apps, for example, simplifies the user experience for native apps like Lync and Outlook, so that the cloud service looks like its hosted in the data center on premise. Similarly, hundreds of other SaaS applications are treated this way. Now the trust model has to be carefully considered. If a company decides to accept federated identities from a separate entity such as Google, then that is a local decision. But I don't see banks ever accepting federated identities from many external organizations, like on-line gaming companies, for instance.

EA: Do you see API security and API access management as growing responsibilities in identity and access management platforms?

PH: Yes, and these might be the most important responsibilities with the expanded use of mobile, tablets, and wearable in the consumer and enterprise environments. Technologies such as HTML5, for example, require that platform solution providers like Ping offer support for the API gateways that result. Even legacy Web access management systems and gateway appliances require support at the API level as they transition to a common enterprise identity and access support model.

Runtime Application Security Monitoring & Protection

Extending Advanced Application Security Controls to the Runtime Operating Environment

Julien Bellanger, CEO and Co Founder of Prevoty

For the longest time, application security implied code scanning, also known as static analysis security testing (SAST), and application scanning, also known as dynamic analysis security testing (DAST). Certainly, the benefits of scanning an application for evidence of vulnerabilities are obvious, and many CISO teams include SAST and DAST in their arsenal. But more recently, the security advantages have become much clearer about embedding runtime controls into the operating environment of an application. So-called Runtime Application Self-Protection (RASP) controls are now emerging as one of the investment areas in enterprise cybersecurity.

EA: What are the benefits of RASP for enterprise applications?

JB: Run-time application security, as we define and employ it at Prevoty, gives enterprise users instant visibility into their production application security posture, not to mention supporting the automatic remediation of existing application vulnerabilities. RASP is unparalleled in its ability to instantly protect your legacy software – that is, those with few if any active developers, while also letting organizations release active applications faster into production, effectively speeding up the secure development lifecycle. Because it can alert on which portion of application code is actually being exploited in production, versus potential vulnerabilities in development, staging, or test environments, development teams can focus on fixing what matters. It makes remediation efforts more targeted and meaningful, saving time and money all around.

EA: If Prevoty's RASP solution runs on production application servers, doesn't that impact the performance and stability of applications?

JB: This question of performance and stability should be one of, if not, the primary considerations CISO teams take into account when looking at Prevoty, or any other RASP solution. After all, the last thing any security program can afford is a tool that negatively impacts applications' performance or stability in production. Through its unique LANGSEC technology, both Prevoty's monitoring and protection capabilities are available with no noticeable impact to the performance of the applications to which Prevoty is attached. We urge readers to explore LANGSEC further with us and understand how this is feasible.

EA: Do you see compliance auditors and regulatory officials becoming more in tune with the benefits of runtime application controls?

JB: Most of our early customers are large financial and commerce enterprises with Web-facing presences and are consequently subject to lots of compliance pressure. Their auditors view RASP as a compensating control for application security risks. We continuously hear regulatory officials asking enterprises, including our customers, to develop and implement actual controls instead of just checking the compliance box. At Prevoty, we've created a product that can integrate with existing vulnerability solutions like dynamic scanners. We've also built integrations with SIEMs that allow auditors and risk management teams to review real-time attack data.

EA: How hard is it for enterprise CISO teams to deploy runtime security? Do they need to fold security libraries into the application code? Or do they run some sort of scaffolding around the application?

JB: In virtually all cases, deployment of Prevoty's RASP in production should be completed in hours, rather than days, weeks or months. There are two models for integrating a RASP solution in large-scale organizations: First, SSDLC code insertion can be performed via SDKs and second, the code can be introduced through plugin attachment via agents. With the former, Prevoty provides support for a wide array of languages such as C#, Go, Java, JavaScript, PHP, Python, Ruby, and Rust, that enable developers to bring security controls directly into their business logic. With the latter, Prevoty empowers Ops and DevOps teams to automatically add security to new and existing applications, leveraging continuous integration and deployment processes for Django, Drupal, Express (node.js), Java, .NET, Rails, etc. Ultimately, we are building an application security product that is addressing complex and large-scale environments.

EA: A big problem in application security has been the weaknesses inherent in the runtime environment such as third party software and components. Do application-

level runtime controls help protect against these weaknesses, or do they undermine the effectiveness of RASP?

JB: If we follow a conservative threat model, we must assume that all third-party software and components, including open source libraries, are vulnerable by default. Furthermore, software that is secure today will become insecure and legacy in the future. By living in the application runtime, such as the Java Virtual Machine (JVM) or Microsoft's Common Language Runtime (CLR), a RASP solution can mitigate against attacks that target vulnerable third-party libraries. For example, our RASP product already mitigated the well-documented Java deserialization attacks affecting many organizations in 2015. For our customers, we were able to save them time, while also reducing exposure and risk.

EA: How well do RASP controls extend to virtual environments? Would the run time controls sit as part of a micro-segment?

JB: Since Prevoty's RASP is attached to an application, it travels wherever the application is deployed: from a local environment, to a physical staging server, to an ephemeral cloud instance. We have many customers today that are moving their monolithic application deployments to micro-services. With this transition, they are using new containerization technologies, like Docker. Today, Prevoty supports applications that run in virtual machines as well as containers.

EA: With SAST, DAST and other existing technologies, CISOs are assured of identifying a broad set of potential application security issues. How does Prevoty's coverage compare?

JB: Recent reports from Verizon and Gartner conclude that over 90% of today's application breaches still exploit SQL injection, cross-site scripting, and cross-site request forgery. So while Prevoty obviously focuses time and attention on these attacks, we are also aggressively improving our coverage model. With our upcoming release, scheduled for GA late summer 2016, we will cover 8 of the OWASP Top 10 categories, in addition to expanded coverage for numerous other attack vectors. All of this is included while our engineers balance the top requirement of ensuring no negative impact on the application's performance.

EA: Where does Prevoty fit within existing enterprise application security programs?

JB: For less mature programs, Prevoty's runtime application security monitoring and protection capabilities can serve as a primary control, providing both detective and preventative measures. For the most mature programs, Prevoty can act as the last line of defense, and can be viewed as additive to a mature program's pen testing, SAST, DAST, WAF, etc. capabilities. There are many gaps Prevoty can fill for those app sec programs between the two ends of the maturity model.

Preventing Email and Social Media-Based Threats

Combining Big Data Analytics With As-a-Service Capabilities Into Advanced Cyber Protections

Gary Steele, CEO of ProofPoint

The reason so many companies try to embed analytic methods into their cyber security protection suite is simple: They work. So when analytics can be integrated into an "as-a-service" offering that focuses on the most likely purveyor of malware into the enterprise environment – namely, *email* – the result is a powerful protection combination for CISO teams. And with business communication moving increasingly to a broad assortment of social media outlets, it is natural to extend these security protections into this new domain.

EA: Would you agree that just about every APT attack we've all seen across our industry over the past few years has involved email-based malware?

GS: Yes, that is unfortunately true. The challenge is that email can be originated anywhere on the planet, and with the openness of email protocols and supporting infrastructure, the objective has always been to deliver such messages to their recipient. This is why email has become such a backbone for global business communication, and also unfortunately the primary mechanism for the delivery of malware. So it should come as no surprise that, as you mention, just about every APT attacks over the past few years has involved email. It's that vulnerability that drives everything we do at Proofpoint to help restore order. And our approach is to combine the best technology for detection and mitigation with support for a comprehensive enterprise approach to reducing cyber risk across the entire range of services in use.

EA: Why do you think the industry has not been quicker to adopt the most advanced cyber security protections for email infrastructure?

GS: I think just about every business and government agency in the world has some sort of email protections in place. The problems, however, are two-fold. First, the attackers have learned to adjust to the most conventional types of security solutions, ones based primarily on signatures such as IP addresses, which are often stale. But perhaps more importantly, most organizations have not developed an integrated architecture for enterprise security with email protection solutions as an embedded component. Our approach at Proofpoint has always had this holistic view at the forefront of every product and service we offer.

EA: What is the value of advanced analytics in a cyber security platform? Do you see this as the secret sauce in detecting zero day and other advanced attacks?

GS: I guess you could say that analytics is our secret sauce, at least in terms of the internal operation of our detection and mitigation platform. It's the underlying algorithms that differentiate one platform from another, and we've demonstrated good success. Keep in mind that for any platform's cyber analytics to be accurate and correct, there needs to be a team of developers consisting of the best and brightest software, protocol, and cyber threat experts – and I'm so proud of our entire team. They've worked hard to develop an industry-leading platform that works well against known and zero-day exploits. But we also know that no team can ever rest, because the most effective solutions today become exactly the sidestepped protections of tomorrow. That's the secret to staying ahead of the offense. The solutions that work today must be reinvented almost as quickly as they are deployed.

EA: How does the cloud impact cyber security solutions?

GS: The cloud is an important part of how we need to think about delivering protection, as cloud-based solutions can update and deploy faster than on-premise tools to stay ahead of the latest advanced threats. For example, recently, our solutions were deployed globally to a 360,000-user organization in just 48 hours. With the ability to quickly deploy and continuously adapt, they allow you to automate the process of detecting, blocking and responding to threats for enhanced protection as your business grows.

EA: So we all know that younger people barely use email in lieu of social media. What's been your experience with social media in the enterprise, and more importantly perhaps, with protecting social media from cyber attacks?

GS: Yes, you are correct. Personal and business communications have certainly expended to include so many more types of services than just email. And yes, younger folks certainly do enjoy social media, although I don't think it's

generational. I think it's more a tendency toward innovation that we see amongst so many different sectors in modern business. Look at how so many experienced marketing teams have gravitated to social media, for example. So we at Proofpoint have been working diligently to apply our solutions to a broader range of over-the-top, messaging, and social communications media. Look at how our researchers detected an infected Android version of the mobile Pokémon GO. We detected that a modified SDK was outfitted to include a malicious remote access tool called Droidjack, which would give the intruder control over a victim's phone. We don't normally think of games like Pokémon GO as being communication media worthy of targeted attack. But that is exactly what we found.

EA: You're an industry veteran. Are you optimistic that enterprise teams will do a better job in the coming years protecting infrastructure? Or do you worry about cyber disasters?

GS: Well, my answer is yes to both of your questions. I do think that enterprise teams will continue to do a better job at protecting infrastructure. There is no doubt about that. We see experts every day in virtually every sector applying the most advanced solutions to tough cyber security problems. Further, ecosystems are a key point of Proofpoint's strategy, and we are partnering with security leaders to build technical integrations bringing greater value for our customers. This year alone, we have partners with Palo Alto Networks, Splunk, CyberArk, and Imperva to give our customers a better security posture when fighting cyber threats. This makes me very optimistic. On the other hand, I also worry about cyber disasters, simply because experience has shown that as offensive attackers develop new techniques, they have the advantage of only having to succeed once, where the defenders have to protect against every possibility.

Providing Visibility to Protect Assets

Enterprise Cyber Security Teams Cannot Secure Assets They Cannot See

Philippe Courtot, Chairman and CEO of Qualys

With computing moving to the cloud, cyber security vendors are scrambling to redesign on-premise solutions. Organizations that don't adopt this new architecture will disappear, like mainframe vendors who ignored the client/server revolution. Cloud computing is also impacting enterprise IT, with CIOs being pressured to move to the cloud to increase agility while reducing cost. CIOs are also realizing that cloud-based security eases deployment and increases scale across global heterogeneous environments. In addition, CIOs see benefit from economies of scale by consolidating on-premise security and compliance solutions natively on the top of a cloud platform. From its inception in 1999 as a cloud-based vulnerability management provider, Qualys had this cloud vision in mind: To provide a continuous view and visibility of security and compliance posture across global IT assets, whether on-premises, on end points, or in the cloud.

EA: Does it surprise you how long it has taken enterprises to move beyond perimeter networks and embrace cloud and virtual environments?

PC: While we saw companies such as Salesforce.com disrupt traditional enterprise software, the security industry has indeed been very slow to enable such a change. This is because ensuring security has a problem of asymmetric nature, where hackers need only to find one vulnerability to penetrate a company's defense, whereas every company must identify all vulnerabilities and embark on the difficult task of eliminating or mitigating them. The result of this imbalance has been a myriad of point solutions to protect enterprise networks. Because of the nature of enterprise software, with its long and complex development cycles, enterprise security and compliance solutions were difficult to deploy and integrate with each

other, which explains the tsunami of data breaches despite major investments we are seeing in cyber security.

EA: How hard is it to locate enterprise assets when they are scattered across hybrid cloud systems?

PC: This is the number one problem, as you cannot secure what you do not know you have, and very few companies know what assets they really have. With a computing environment where complexity and scale are increasing; where almost everything connects with everything else; and where the data is scattered across many different environments, a new architecture is required to continuously identify, catalog and assign attributes to companies' global IT assets. This is becoming the number one priority.

EA: How important is continuous monitoring in the process of vulnerability management?

PC: It is critical to have an updated inventory of all corporate assets and to continuously identify those out of compliance, either because they have vulnerabilities or misconfigurations that can be exploited, or because they are in violation of internal policies or external regulations. In fact, IT, security and compliance should be brought under a single integrated solution, rather than having three different cyber security solutions that are in fact looking at the same data in three different ways, with no way to correlate the results. You simply cannot effectively protect what you do not know is there.

EA: Will the growing enterprise use of data encryption change the vulnerability management process? Will scanning tools, for example, be able to find and actually see what they are looking for?

PC: Indeed, encryption is going to become more pervasive as the industry moves toward protecting data both at rest and in transit. As a result, security must be built into the fabric of today's cloud computing environment, rather than through retrofit security afterwards. This proactive approach can be accomplished by embedding sensors, which we call cloud agents, into every component to identify vulnerabilities and unusual behaviors while looking at the computing environment like hackers do by using dynamic and passive scanning. With such an approach, we can look continuously at this new computing environment from both the inside and the outside to detect potential vulnerabilities or malware that could result in breaches and to be able to automate the detection and response to attacks.

EA: Philippe, you've had such a broad perspective across the technology and security communities during your amazing career. Does it seem like the security community faces more of a challenge today from the rapidly changing world of business and IT, or from rapidly advancing methods of hackers and criminals?

PC: As I have had the opportunity to explain many times to the security community, these rapid changes becoming evident today are just like previous disruptive changes in the history of computing. They present both a unique opportunity and a formidable challenge. Simply said, those who do not embrace the change will be left behind, as nobody can resist the power of creative disruption for long. So, welcome to the cloud-computing era.

Securely Delivering Important and Sensitive Documents

Secure Delivery of Important and Sensitive Records Requires More Than Just Data Encryption

Mark Morley, COO of SecuritiNet

The wall of paper folders and records displayed behind the receptionist's desk at your doctor's office is a good metaphor for how well industries such as health care have adapted to the proper use of modern technology. The good news is that hackers probably cannot access paper stuffed in a folder. The bad news, however, is that as the medical, health, insurance, and related industries have begun moving to more electronic means, including mobility and cloud, where standard security solutions are not often found, the risk of compromise increases dramatically. Securing the delivery of such important and sensitive information in this context and other large sectors is becoming a massive priority in virtually every sector.

EA: Why do you think that so much important information is transmitted so insecurely?

MM: First of all, few would argue that email continues to be the main source of business communication, even for the transfer multi-million dollar merger documents. Since email does not include much security by default, such transfer is inherently insecure. Second, much of what people refer to as "work" is really just a progression of ad hoc tasks, often short-term, that may not have been designed to include secure transmission capability.

EA: Are companies starting to use encryption more regularly?

MM: They are trying to use encryption, but the tools and set up are generally so hard to deal with, that most people just give up. As an illustration, a friend of mine with no less than *three* degrees in Computer Science, including a Ph.D., recently tried to encrypt and send some business correspondence to a colleague. He told me that

he tried multiple methods for nearly an hour and then gave up. I think he decided to just put the letter in Dropbox. My observation is that this is pretty typical today.

EA: Are there government regulations in the United States or other countries governing the risk of information sharing? Maybe this can help keep important correspondence out of plaintext cloud folders.

MM: Yes, there are reasonably mature standards in the payment card and health care industries. Additional standards for secure file sharing are starting to proliferate in other industries. While these standards are mostly advisory, this is certainly headed in the right direction. The bottom line is that we need technology to make it easier for people to do the right thing when sending important and sensitive documents. It is far better to have an easy-to use, highly secure system instantly available, than to rely on your employees to always make good decisions.

EA: Why do you suppose that stealing medical information is so popular now? Is it true that personal medical records are more valuable than credit card information?

MM: A stolen credit card is sold for cents, whereas a set of Medicare ID numbers can get you several thousand dollars. Incidentally, five thousand dollars is worth about 22 bitcoins, in case you were wondering. Criminals know that health care records have a longer life than credit card numbers, which can be cancelled quickly. And the possibilities for committing medical insurance fraud are endless.

EA: You mentioned the health care and the credit card industries a couple of times now. Are they unique in their data and file exchange requirements? Or are all other industries just as challenged?

MM: Other industries are definitely challenged. The high percentage of work that is basically unstructured and done by email, can be hacked by high-speed robots that can read all this material at line speed. For example, read your Google license; it says that Google is going to read what you send – and this applies to every industrial sector and consumer application. So, if Google can read it, then the bad guys can certainly read it, since email is routed across the public Internet. After looting the high value documents, the bad guys will move down to the medium-value documents. That's the path we are on, unless industry sectors begin to recognize the value of using new networks, new technology, and improved secure file sharing techniques.

Securing and Enabling Cloud Services for the Enterprise

Cloud Access Security Brokers Protect Corporate Data Across Public and Hybrid Services

Rajiv Gupta, CEO of Skyhigh Networks

When employees rush to public clouds for service, there is usually good reason. And while many CIOs hesitate to admit this, the capabilities offered in cloud services such as Office 365, Box, ServiceNow, and Google Drive often exceed the capabilities provided by the local IT staff. Ultimately, this is good news for CIOs once they accept and embrace the notion of infrastructure, services, and computing being provided in an "as-a-service" manner. The biggest hurdle, however, involves assurance that security, compliance and data protection are properly handled. This is where cloud access security brokers (CASBs) have emerged as an important component in modern architectures.

EA: First of all, do you think the traditional notion of a perimeter works to stop cyber attacks any more?

RG: The reality in today's workplace is that the perimeter is rapidly disappearing. Employees are working remotely from home or from their favorite coffee shop, and accessing corporate systems from their personal devices. Partners collaborating from outside the perimeter are securely sharing confidential data. Therefore, traditional perimeter-based security needs to be re-imagined for the cloud-first and mobile-first era.

EA: CIOs have referred to public cloud usage in the past as part of shadow IT. Do you think this perception is changing and do you think the time will come when the majority of IT functions are delivered from the cloud?

RG: Over the last four years, we have seen a major shift in CIOs' attitudes towards public cloud services. There is universal acknowledgment that public cloud services can help accelerate innovation, increase employee productivity, and drive the business forward. As trusted partners to their business counterparts, CIOs have shifted focus on enabling cloud services, while meeting their various security, compliance and governance requirements. As CIOs embrace cloud services aggressively, I predict that one day a majority of IT functions indeed will be delivered from the cloud for a broader number of companies than today where today this largely is true for younger companies.

EA: Have large and small companies behaved differently in the adoption of public cloud services?

RG: In the early days, IT departments at smaller companies in non-regulated industries led the charge to the cloud. This was partially due to necessity (they didn't have the IT staff to support on-premises software), and partially because they had less stringent security and compliance requirements. IT teams in larger companies carefully plotted their adoption of public cloud services. However, business units and employees in the larger companies often lost patience and, taking a page from nimble smaller companies, took the initiative to adopt public cloud services, resulting in the shadow IT phenomenon.

EA: Tell me how a cloud access security broker function works. Is this a man-in-the-middle solution?

RG: A Cloud Access Security Broker is a cloud control point, one that controls access to cloud services, and protects corporate data in cloud services. Full and comprehensive CASB functionality – which includes visibility, threat protection, compliance, and data security – requires that the CASB support all deployment modes, including log-based visibility, integration using proxy and firewalls APIs, forward proxy-based inline intermediation, reverse proxy-based inline intermediation, and integration with APIs provided by cloud service providers. Different customer use cases require different deployment modes. Log-based visibility and integration with cloud service provider APIs are off-line approaches. Offline approaches can't address requirements such as data encryption, data jurisdiction controls, or access control based on context which inline controls can. Integration with proxy and firewall APIs leverage existing firewall and proxy solutions to intermediate cloud data access. Forward and reverse proxy-based intermediation approaches are inline to the cloud access, however there is a big difference: reverse proxy-based intermediation does not require any device agent or other footprint, while forward proxy-based intermediation does. In traditional parlance, forward proxy approaches are referred to as man-in-the-middle.

EA: For complex environments with public, private, hybrid, and even traditional enterprise IT, is it a tough project to design, integrate, install, and operate a cloud access security function?

RG: In general, the complexity and effort of deploying and maintaining a cloud access security function depends on the deployment mode chosen. In general, off-line CASB deployment modes are faster and easier to design, integrate, install, and operate. Reverse proxy-based approaches that integrate with Single-Sign On solutions are similarly fast and easy. Forward proxy-based approaches are generally the most challenging, because they require device footprints to be distributed, installed, and maintained on all devices even as device versions change over time. Our experience with some of the largest organizations in the world, including 25% of the Fortune 100, has taught us that immediate and significant ROI can be delivered first through a cloud-based product that is simple to configure and integrate with existing in-cloud on on-premises IT investments; second, through best-practices workflow derived by hundreds of customers and codified in the product; and third, through a 5-step deployment methodology that delivers greater value over time. As a result, Skyhigh customers begin to see immediate value in as little as a week, and it grows from there.

Hunting Down Cyber Attacks in Enterprises with Big Data

Threat Hunting Platforms Support Advanced Analytics to Detect Cyber Attacks in the Enterprise

Adam Fuchs, CTO of Sqrrl

A promising shift in enterprise cyber security is the trend toward proactive hunting of threat issues in advance of their causing consequential damage. Previously, cyber security analysis consisted of collecting data from gateway systems that would passively watch as an attack occurred. This collected data would be passed to analysts who hopefully would recognize what was happening in order to initiate response. By shifting this approach to a more proactive approach offers hope that attacks can be stopped before they are completed.

EA: Is security analytics anything more than just correlating collected data?

AF: Absolutely. Security analytics is the application of advanced algorithms, including supervised and unsupervised machine learning and graph algorithms to identify threats that evaded detection (or proper prioritization) by other security systems. Many of these algorithmic techniques have been around for a while, but one of the major advances for security analytics today is the ability to deploy them at scale across massive amounts of data. However, it is not enough to just correlate and automate analysis at scale. Textbook application of machine learning techniques frequently produces groundbreaking insights, like "http traffic is often seen on port 80," thus leaving a fair amount to be desired. To truly impact the security domain, analytics require structure and context; structure in the form of behavior and attack models; and context in the form of broader prospective from multiple sensors, risk analysis, and feedback. With appropriate structure and context security analytics are an incredibly valuable tool for hunting, detection, and forensics.

EA: How hard is it for security analysts to learn to hunt attacks? Do analysts need to be experts in networking, mathematics, and investigative forensics?

AF: Sqrrl is the provider of the leading threat hunting platform for security operations centers (SOCs), and our assumption is that our customers don't need any data science skillsets. Historically, to proactively hunt for threats you needed to have data science skillsets to build custom algorithms to look for anomalies that other tools missed. Sqrrl's algorithms work out-of-the-box on standard datasets seen in most SOCs. The structure for Sqrrl's analytics comes from extensive modeling contributed by Sqrrl's security experts and experts in our network. Sqrrl deployments learn much of the necessary context through observing data feeds from a variety of sensors. From the start, Sqrrl's analytics find interesting behaviors that give insights into what's happening on the network. Analysts provide feedback on false and true positives over time to hone in on exactly the behaviors that matter to them.

EA: How does the enterprise transition to mobile devices and cloud systems affect the security analytics process?

AF: One of the big changes with increased use of mobile and cloud is that enterprises are starting to give up on the idea of a secure perimeter. With attack vectors in email, web browsing, and countless other common activities, secure perimeters have been a dubious concept at best for over a decade. Mobile and cloud systems are acting as a forcing function for companies to break old habits and begin adopting more effective tools and techniques. For those of us already in the modern world, the biggest change we see is in our sensors. Mobile and cloud systems make some behaviors harder to spot and other behaviors easier. With well-structured security analytics, we can take advantage of new datasets that can provide additional detail and context into potential attack pathways and attacker TTPs (Tactics, Techniques, and Procedures). As an example, security analytics tools can take logs from Cloud Access Security Brokers (CASBs), and correlate behaviors associated with them to look for data exfiltration patterns and connect those patterns to other TTPs correlated with the same hosts and users.

EA: What sort of trends do you see in cybersecurity vulnerabilities in the enterprise?

AF: In general, we are seeing increased cyber security awareness and better cyber hygiene in large enterprises. However, many attacks do not require exploitation of traditional software vulnerabilities. These exploit-less attacks often take advantage of human vulnerabilities and then move laterally and escalate privileges without the use of malware. This is why enterprises cannot rely on anti-malware or anti-virus solutions as a sole layer of defense.

EA: Do you think that smaller companies can ever take advantage of security analytics tools directly? Or do they need to rely on managed security service providers with trained staff?

AF: We believe it is critical to still have a human in the loop when conducting threat hunting. Fully automated solutions can only get you so far. As a result, we do see benefits in smaller companies taking advantage of the MSSPs. Recruiting, training, and retaining advanced security personnel are difficult for larger companies, let alone smaller companies. MSSPs can help mitigate this, and more and more MSSPs are now offering specialized threat hunting services to their customers. Beyond just the expertise consideration, MSSPs are also uniquely positioned to correlate attack indicators across multiple companies. This really helps to identify signals in the noise and pick out potential attacks earlier.

EA: Do you think it is realistic for an enterprise to ever hope to detect attacks from advanced nation state actors? It seems like an unfair fight.

AF: No organization can guarantee that a well resourced, determined adversary will not be able to breach their perimeter security. However, threat hunting and security analytics can greatly assist enterprises in reducing the probability that such an attack will be successful. Sqrrl has assisted a variety of Fortune 2000 companies, government agencies, and MSSPs in detecting these types of advanced nation state actors.

Access and Analytics to Support Cyber Investigations

Using Network Data to Support Cyber Security Analytics and Advanced Case Investigations

Faizel Lakhani, President and COO of SS8

Collecting data on a network at line speed for the purpose of analysis, management, and intelligence is one of the basic tenets of how telecommunications services have always been provided. Law enforcement relies heavily on the use of these techniques to make society safer and to embed strong disincentives for anyone to break the law using communications networks. Now, as data breaches continue to occur despite preventative security, we are seeing these practices applied to the enterprise for more proactive data breach detection and response.

EA: Is it getting harder to extract the right data from a network to derive useful intelligence?

FL: It certainly is harder. In the early days, networks were much more predictable with point-to-point connections and hub-and-spoke architectures. It was easier to understand where traffic was going and which applications were being used, because the apps were centrally hosted in the data center. Fast-forward to today and everything has shifted to the cloud. There are now countless ways to share information, and this is giving nefarious people countless ways to exfiltrate proprietary information. The use of encryption has confounded much of the analysis of communications resulting in opaque and hard to understand flows over a network. While all of these combined factors have made the challenge of connecting the dots for a cyber investigation or for breach detection exponentially more difficult, we believe this is an area where we've really "cracked the nut" with our Protocol Extraction Engine, or what we call PXE (pronounced "pixie"). PXE is a powerful, highly optimized deep packet inspection (DPI) engine that is capable of classifying more than 1,000 protocols and performing metadata extraction for several hundred. The High-Definition Records (HDRs) we generate from the

network go beyond weak indicators like port numbers to classify protocols. Instead, SS8 uses behavioral DPI including packet flow analytics to classify traffic by inspecting flows, which succeeds even if tunneling or obfuscation techniques are used.

EA: Do you think there are privacy-related showstoppers that will make it impossible for law enforcement or to gain the information it needs to catch bad guys on the Internet?

FL: Not impossible, but certainly more difficult due to a focus on encryption. It's a fine line we must walk to maintain doing what's in the best interest of detecting versus observing. In the construct of enterprises, the use of encryption makes it hard to use traditional methods of detection to find and remediate breaches. This is where our years of history helping law enforcement and intelligence agencies has enabled us to build competency in understanding and mining encrypting communications to extract information that can guide on the intent of the communications. The struggle for law enforcement and intelligence is that any events of criminal activity or terrorism are based on coordination, and this coordination is over communications. Hence when this communication goes increasingly dark, the challenges for technology become much higher. This is where we built our competency.

EA: Virtualized, software-defined networks seem to be on the rise. Are these networks making it more difficult to detect malicious activity?

FL: Virtualization has changed the game a bit, and today, is a bit of a blind spot for the enterprise. You can no longer rely on layer 3 and 4 network analysis tools such as NetFlow to understand a full picture of what's happening. Advanced threats will jump between workloads to hide in the normal flow of communications, providing a great risk for data exfiltration. Organizations need to be inspecting traffic at the application layer and all the way up the stack. This requires running the communications analytics sensors as virtual instances. And rather than having to find the natural egress points across the network, SDN affords the opportunity to perform data collection in one place.

EA: As just about everyone shifts to mobility-accessible public cloud applications, does this change the nature of data collection and analysis?

FL: The explosion of public cloud applications in recent years has made the collection and analysis of data for cyber investigations more challenging. Everything is going over the Internet, and nearly everything is encrypted. We must now look at extracting meaningful summarization for applications running on the network, and continuously build coverage for popular mobile applications. It is this constant motion that requires SS8 to dedicate a team to understanding, mapping and decoding the packets from these communication applications. And in light of ever-

increasing use of encryption, we have built advanced features that reveal useful high definition records about certain types of encrypted sessions. For example, the encrypted call detection feature in our software is able to pinpoint and differentiate voice and video calls over encrypted services such as Skype and Viber. Another example is with data breach investigations, where we can look at the certificate data and uncover when the signing authority does not match the source. Encryption is a challenge for deep packet inspection and developing techniques to provide useful metadata in spite of encryption is a critical part of SS8's feature set.

EA: Do you see security analytics moving in the direction of more automated solutions, rather than as toolkits for human analysts?

FL: Absolutely. Automation is essential across both cyber investigations for law enforcement, and for data breach investigations conducted by enterprise security analysts. There is no slow-down in the amount of threat intelligence coming in, but there is however, a shortage of experienced cybersecurity analysts who understand how to process the information or to even understand if it is relevant. The process of intelligence-to-action is manual today, and the goal of SS8 is to automate that. It's the difference between ingesting new threat intelligence and manually having to connect the dots. Automation ensure that the dots are connected for you to pinpoint an unknown suspect or device of interest for the analyst to then determine if the dots paint a worrisome picture of a random collection. The next generation of software can take in the latest threat intelligence automatically and constantly match it against the history of high definition records from the network. As new learnings and threat discoveries occur, alarms go off about an indicator of compromise, and the suspect or device of interest is identified. As a secondary level, similar to seeing a known bad actor on the street, this new model of breach detection will look for further actions that make it obvious that someone needs to be called to investigate. The human element can't go away completely, but it's about allowing those security analysts and cyber investigators to do more with less.

EA: You seem to be moving toward enterprise breach detection. Is there any correlation between how you collect information for law enforcement and for the enterprise?

FL: It's no secret how problematic breach detection is for today's enterprise. It seems like there is a new headline each day about a company being breached. We've all seen the numbers: breaches going undetected for more than 240 days, with most breached companies finding out they were breached from outside their organization. Our deep understanding of communication flows and years of proven experience tracking suspects-of-interest (SOI) has given us a unique edge in being able to rewind and pinpoint the device of interest in today's war on enterprise data breaches. It's all about taking today's knowledge and applying it to history. Our breakthrough Learning Analytics model ties together the high-definition records from communication patterns with today's threat intelligence to not only accelerate

the detection of breaches, but to forecast breach behavior for future protection. Network history offers the fastest means to uncover the unknown, and only when you constantly wind the clock back using the latest threat intelligence and network history can you uncover what gets missed by preventative security.

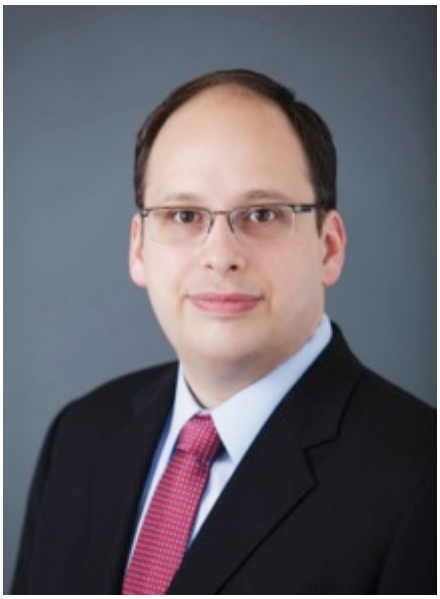

Integrating Endpoint, Network, and Cloud Into a Secure Ecosystem

Combining Proxy-Based Inspection, Threat Intelligence, and Web App Protection into an Enterprise Solution

Dr. Hugh Thompson, CTO of Symantec

Enterprise security has been well positioned to take advantage of proxy-based traffic inspection and policy enforcement. The natural chokepoint that exists for every company's Internet and other external access has provided a convenient means for enforcing URL, application, and other security policies. The unique architecture of a proxy enables a deeper level of control over communications and inspection of the content not available in more simplified approaches. Now that the enterprise is virtualizing around cloud and mobile technologies, this important monitoring and mediation function must shift toward integration with less perimeter-oriented LAN approaches.

EA: Every enterprise is more dependent on their wired and wireless networks than ever before. How do you see the corporate network evolving in the next few years?

HT: Network connectivity is just fundamental now. We've all witnessed the increase in the mobility of workers, exemplified by the use of WIFI within offices, coffee shops, and at home. And the usage of 4G/LTE networks – which results in connectivity *everywhere* – will only accelerate more with 5G wireless services and the associated increases in data capacity and speed. In this new environment, employees will untether from corporate WIFI more often, and billions of new devices will be constantly connecting. As network access continues to evolve and improve, deeper network security will be more tightly bound to that access onramp.

EA: Are you seeing an acceleration of migration to cloud? Larger businesses, especially the banks, have seemed a bit more cautious. What's been your experience here recently?

HT: The caution from certain sectors is understandable if you consider the compliance and regulatory requirements that many of these organizations have to deal with. We've often seen companies start by working to gain an understanding of cloud usage and risk within their organization – that is, understanding the non-sanctioned cloud use by individual employees or departments. The conversation then quickly turns to controlling that unsanctioned cloud – usually by blocking the riskiest applications, and by forcing inspection by DLP and threat protection tools. All of this is motivated by the need to mitigate compliance risks. And the challenges are even greater on the cloud apps that organizations are sanctioning. Additionally, we've seen data residency issues from the European GDPR and Asian data sovereignty laws impacting global rollouts of large cloud initiatives. We've worked with CISOs trying to understand how to integrate cloud activity into their incident response processes. We've worked to highlight the potential compliance risk of content and permissions in cloud applications where we've seen how our content inspection and sandboxing capabilities can identify advanced threats stored in cloud content or transiting via email. Overall, enterprises are quickly learning what needs to be done to extend governance and security processes to cloud. That's why we've aggressively acquired and built solutions into our platform to help them achieve their goals and bring a defense-in-depth approach to securing the transition to cloud applications and services.

EA: Will it be easy to virtualize proxy-based security? Every enterprise architect in the world starts by sketching the proxy into the perimeter. Where does it go with cloud-based architectures?

HT: Proxies are already virtualized today, as well as being delivered as a cloud service, and there are some important distinctions. Virtualization allows organizations to deploy proxies in any part of a private cloud or in IaaS architecture where they want to create a security control point, usually as a micro-perimeter. Furthermore, proxies as a cloud service allow any type of device, at any location, to gain that proxy-based protection, making it that much easier. Enterprises use a proxy cloud service to enable safe branch office access to the Internet, protecting mobile devices that have no agent-based solution. We've seen use by automobile manufacturers to protect Internet-connected vehicles, demonstrating a security model for consumer devices, industrial controls, and the next generation devices often referred to as the Internet of Things (IoT). Virtualized and cloud-delivered proxies make it easier to architect that protection into any device or application model.

EA: How do you see the cloud access security broker playing in enterprise security? Is this the new perimeter?

HT: I think we've already moved away from a concept of a single perimeter. Cloud access security broker capabilities are critical components to extending security and governance to cloud applications, but it's an additive problem. Cloud apps present

whole new sets of use cases related to many different areas of IT and security. These include discovery and access control of shadow IT; activity logging for compliance, breach detection, forensics, and incident response; data compliance issues around data residency laws; access rights for documents stored in the cloud; protection of documents with PII, PHI and other compliance sensitive information; and granular policy enforcement to prevent theft of data and identity. But, even as some applications, data, and content have shifted to the cloud, others remain within the enterprise domain. Enterprises need to figure out how to solve for the new problems, and make the solution consistent, integrating it with their existing security and governance infrastructure.

EA: With so much complexity in modern networks, how does a CISO team find ways to simplify? Everything seems so complicated nowadays.

HT: The challenge of the CISO has definitely gotten more complicated. It's a tough job, calling for the best people to navigate problems, technology, processes, and partners. The best approaches we've seen start with an immediate need – a pain point, so to speak – but to also draw in strategic and long term considerations. CISOs understand that if their teams spend all their time stitching together narrowly focused technologies from different vendors, they'll never achieve the operational effectiveness that they need. The smartest CISO's choose their partners strategically, but also insist that those partners and platforms remain flexible and open to accommodating changes in requirements. In our experience, this seems like the optimal approach.

EA: Any thoughts on whether the defense ever catches up the offense in the cyber security game?

HT: We all sure hope for that to happen, but we must plan differently. We have no choice but to expect that the criminal offense will continue to advance and evolve. And, of course, the defense will also continue to innovate to improve security. For that reason, Symantec fundamentally believes – as has always been the case with the Blue Coat team – that the best approach to improving one's defense is to maintain an open architecture. This allows enterprise security teams to quickly integrate innovation to improve their defense posture, and it simplifies the ongoing operation that streamlines security and governance processes. The proxy serves a critical role to enable Web and cloud governance, threat protection, data protection and incident response. We view our open platform architecture as a way for our customers to continuously extend and adapt their security posture in the face of evolving threats.

Using Trusted Hackers to Reduce Security Risk and Management Chaos

How a Global Army of Ethical Hackers Can Be a Trusted Part of Your Everyday Enterprise Security Team

Jay Kaplan, CEO of Synack

Historically, when security researchers, or ethical hackers, have discovered security vulnerabilities in an enterprise's applications, infrastructure, and products, it was unclear how and where they could disclose this information in a safe, legal manner without facing reprimand for their actions. Even in the absence of malicious intent, ethical hackers may face legal backlash from companies who have been hesitant to implement vulnerability disclosure or bug bounty-like programs that invite external researchers or hackers to submit discovered vulnerabilities discovered across their digital attack surface. Recent advances in private crowdsourced vulnerability disclosure programs can facilitate and even reward hackers for helping find to fix security vulnerabilities that have gone undetected by more traditional security measures. Today, some of the largest organizations in the world from both the public and private sectors are learning that with the right structure and control, leveraging a broad base of security experts that are incentivized to discover vulnerabilities may be one of the most effective measure to combat the onslaught of attacks they're facing today.

EA: Is it accurate to describe Synack's offering as a bug bounty program?

JK: Actually, we like to describe our Synack solution as a Crowd Security Intelligence Solution. Our model shares the crowdsourced, incentive-driven, and adversarial components of most bug bounty programs. But with the complexities of modern cyber attacks and the difficulties in identifying and managing the vulnerabilities that can exist in today's enterprise infrastructure, we've designed our approach to be more holistic than the typical bug bounty program. We've tried to develop a full-service model centered on the concept of trust and ease of use.

EA: When you say trust, are you implying trust across a community? And do you mean a public or private community?

JK: As I mentioned above, unlike the typical *public* bug bounty programs, Synack cultivates a diverse and *private* community of highly curated security researchers. We refer to this group of experts as the Synack Red Team (SRT), all of whom have been carefully vetted for both skill and trust, and their activity is continuously captured and monitored through our LaunchPoint platform, a characteristic that most existing bug bounty programs lack. This combination of vetted security researchers, combined with an auditable activity record, provides full transparency and sufficient technical controls for even the most conservative organizations to take advantage of crowdsourced application and asset testing for sensitive applications and internal environments.

EA: For Synack customers, are you scanning target environments, perhaps at Internet visible entry points?

JK: In addition to the focus from our SRT, all Synack engagements will benefit from Hydra, our proprietary technology that continuously probes and scans the assets and applications in scope. This approach enables our researcher community to more efficiently scale their testing and vulnerability discovery activities, and is better situated than traditional bug bounty programs to meet the needs of clients who manage vast and rapidly evolving collections of assets. When executed by large corporations who have no problem attracting (and affording) top security professionals and can allocate the appropriate resources to efficiently manage, triage and support bug bounty programs – they can surely be effective. However, without all of the necessary resources, the excessive noise and lack of accountability and trust of bug bounty programs can become problematic and overwhelm internal security teams.

EA: What happens if one of your vetted researchers identifies and reports vulnerability that is already known? Does the researcher get compensated?

JK: No. A vulnerability that's already known by the client organization is called a duplicate submission. Synack, and most incentive-based programs, typically reward only the first participant to report a discovered vulnerability. This motivates researchers to be constantly on the hunt for vulnerabilities, and to report discovered bugs in a timely fashion, before someone else reaps the reward. A well-run program, like Synack's, will let the hackers know what vulnerabilities are known ahead of time and therefore out of scope, so as to respect the researcher's time. In rare instances, duplicate submissions may be rewarded when taking into consideration various factors such as the severity of the vulnerability discovered and quality of the report submission and accompanying details.

EA: Have there been business obstacles for enterprises to begin adopting your Crowd Security Intelligence solution?

JK: The adoption has been incredibly strong. Initially some CISOs do struggle with the concept of having hackers attack their digital assets, but they get over these objections, usually when they come to understand the process and controls that are in place. We explain the careful reviews, screening, and testing that are in place for all Synack Red Team procedures, and we elaborate upon the assurance and accountability that is achieved by having all SRT activities tracked through Synack's proprietary LaunchPoint technology. Additionally, we make certain that they understand the absolute confidentiality that is provided to all customers on their identity and any discovered vulnerabilities. These attributes and controls typically ease business concerns and remove obstacles for security teams to begin working with us.

EA: Have you seen a change in the quantity and quality of reported vulnerabilities since bug bounties have been in place?

JK: As Synack's client base and researcher community has grown over the past three years, the volume of vulnerability submissions has followed suit, but our commitment to quality has always remained our top priority. We pay close attention, as our customers do, to the signal-to-noise ratios (SNR) in our reporting. Our internal Synack Mission Ops team provides comprehensive vulnerability triaging, validation, prioritization, and reporting, which has resulted in a SNR of over 95% across all Synack engagements. This means that our customers prioritize over 95% of the SRT-submitted vulnerability reports they receive from Mission Ops as "must-fix" vulnerabilities, with less than 5% of reports they receive being categorized as duplicate, out-of-scope, or "won't fix". We emphasize valid and actionable results over pure volume returns to avoid focus on non-exploitable vulnerabilities and duplicates that can overwhelm a security team, wasting their valuable time and resources.

EA: What are the prospects of bug bounty services for SMB? Do you think subscriptions can reach that size company?

JK: It's inevitable that crowdsourced security testing programs will reach the small business market. We've already seen it with some of the small, but quickly growing tech companies that have adopted Synack as part of their security lifestyle. Nowadays virtually every company, independent of size, is a data company and has layers of technology built into their business functionality and day-to-day operations. So it is naïve to think that only large companies with sophisticated security programs can embrace our Crowd Security Intelligence solution. Organizations of all sizes are witnessing the value and benefits of Synack's crowdsourced penetration testing and continuous application security testing over more traditional application and network security measures. Our fully-managed,

cloud-based solution that harnesses the skills and expertise of hundreds of the best security researchers from across the globe is the most effective way to scale security testing across dynamically changing and expanding attack surfaces, and eliminate hiring and vendor onboarding delays that prevent on-demand testing engagements. Additionally, as we continue to evolve Hydra, our proprietary automated scanning technology, we will make continuous Crowd Security Intelligence testing available in a cost effective way for mid-market and SMB customers.

EA: What's been the weirdest thing reported that you've seen?

JK: Per confidentiality agreements with all Synack clients, we don't publicly disclose specific vulnerabilities discovered across our diverse customer base. With that said, recently a Synack Red Team (SRT) member reported an extremely severe vulnerability in a critical enterprise system that required users to enter their Account Number and PIN to authenticate access. If you provided the wrong PIN, it would block access, but if you inputted no PIN whatsoever and pressed submit, it would grant access to any Account Number.

Supporting Programmable Biometric Authentication

Going from Requirements to Implementation Without Additional Software Development

Rakesh Loonkar, President of Transmit Security

Sadly, passwords remain the primary means by which most on-line banking and retail services authenticate their customers. This is not surprising, given the resilience of passwords, their interoperability with so many different systems, and their perceived low cost. But advances in the accuracy and dependability of biometrics, combined with improvements in their underlying infrastructure support in the enterprise, make biometrics a more viable option for so many new and existing authentication needs.

EA: Should enterprise security teams start the process of getting rid of passwords from their applications and systems?

RL: I think that total removal of passwords from an enterprise may be an unrealistic goal, but everyone knows that biometrics can be more secure, and a better option for many types of business and security requirements. At Transmit Security, we've tried to make this transition simple and easy for technologies such as eye recognition, voice recognition, facial recognition, fingerprint, and all different flavors of one-time passwords. Years ago, these all seemed like technologies appropriate for use as NASA or a high-security government agency. But today they can be integrated into virtually any IT environment – and that is exciting.

EA: Have biometric factors gotten to the point where they are sufficiently accurate to serve as truly trusted means for accessing bank accounts and other important assets?

RL: The biometric accuracy is not the issue – this technology works in a dependable and secure manner. The problem is that when passwords and other weaker forms of authentication are used to enroll fingerprint users. Yes, this enables the convenience

of using biometrics, and that might be just fine for some applications. But in places where the security is paramount, Transmit's solution helps ensure end-to-end secure authentication, including registration, for users.

EA: How does adaptive authentication work?

RL: First, "adaptive authentication" is a very generic term and it historically has been applied to the combination of an authenticator with a device ID. Enterprises have an issue implementing truly adaptive authentication for two reasons. First, it's very difficult for enterprises to hard code every exception and orchestration use case into their applications, and second, as new modalities of context become available on the market, the internal costs to acquire, integrate, and orchestrate the output of these tools are very high – typically in the millions of dollars for a large enterprise. When enterprises think of *adaptive authentication*, they should be thinking in terms of how to eliminate most of the internal software development steps, be able to add new context, and change an authentication process in any live application, for any reason, in a matter of minutes or hours. That would be truly changing the game.

EA: What does it mean for an authentication system to be programmable?

RL: Programmable means that you can go from requirements to implementation in minutes without re-coding your applications. That means that you can go from delivery times of months or years for identity related projects to a matter of minutes. Programmable authentication, as we have implemented at Transmit Security, allows for an enterprise to turn any biometrics, device level authenticators, context, anti-fraud indicators, or tools and build and change logic around them. Programmability allows security and IT teams to basically off-load the work of embedding authenticators into every application, which can be a monumental task. Instead, all of this logic can be abstracted out of the application and make use of a common interface, or API, for authentication and provisioning tasks. Furthermore, programmability supports selection of authenticators, and anti-fraud tools based on the needs of users and the business. This allows for application owners to literally select their desired approach, whether it be facial, voice, eye, OTP, push notification, or some other means, or any other authentication process. Adaptive context can then be added and combined with behavioral profiling to create a super-secure environment with an awesome user experience.

EA: How does the analytic process integrate with authentication? Does this require connectors between tools like the SIEM and the authentication platform?

RL: You certainly could connect these different platforms, and we have many customers that do. But the Transmit Platform includes native support for behavioral learning. Users simply provide information about profile targets, and the solution automatically combines information about devices, access times, location,

transactions, and other factors. The idea is that each user would be associated with a profile based on their behavior, and this would allow business owners to keep thing secure, and keep users happy. The alternative for the enterprise would be to procure multiple systems, integrate themselves. This typically leads to several multiples of higher costs and horrible delivery times of even basic functionality.

EA: If a company already has an authentication solution in place, do they generally have to toss the whole thing to make improvements? Or can they do something incrementally to increase strength and options for users?

RL: That's a good question, but in general, enterprise security and IT teams rarely have to toss entire systems to support better authentication. The Transmit platform is designed with a simple interface that works with new and existing systems in order to avoid precisely that situation. No team likes to remove systems that have been invested in, and nurtured. The better option is to integrate, and we certainly can do this for new and existing identity and access management, application, system, and network infrastructure. And this also goes for security analysis solutions and processes. Everything should work together. However, we are finding that over time, Enterprises want to consolidate their many systems because they are priced at a premium, but perform commodity functions. So we see customers implement our platform for one function, but then use more of the platform functionality as they see the value.

EA: You referred to your solution as being omni-channel in certain marketing literature. What do you mean by that?

RL: We're supporting all of the mobile device level authenticators that everyone is thinking about and allow any enterprise to connect them, with risk processing in the middle, to their channel applications. Examples of these applications could be the call center (such as IVR), the branch or store, or an ATM or Kiosk, and of course the Web and mobile applications they own. Enterprises want to be able to offer their users a unified user experience for identity across all of these channels. The Transmit platform, after a relatively easy one-time integration, allows any enterprise to implement thousands of use cases in any specific channel and across many channels. For example, if you want to use iBeacon in a branch, then no problem. If you want to use eye recognition to verify a user for call center verification, then no problem. I can get more sophisticated here, but you get the point. All of these use cases can be engineered and implemented in a matter of minutes.

Supporting Anonymous Cyber Incident Exchange and Collaboration

Providing a Trusted Network for Industry to Share Incident Anonymously

Paul Kurtz, CEO of TruSTAR

Information sharing and collaboration are topics in cyber security with universal appeal, but also with relatively weak adoption to date. Many reasons exist for such slow progress: First, the basic mechanisms to support cyber information sharing have been slow in development. Second, the laws in most countries governing how information can and cannot be shared are often fuzzy. Third, the attribution associated with more specific incident information can leave an organization at risk. And fourth, there has been little actual *return* on cyber information sharing for participants.

EA: Everyone agrees that information sharing is important – so why do you think it has received such weak adoption?

PK: The gold standard of incident sharing, the one that everyone wants to be a part of, is one where active incidents are shared broadly. Buried within this is the expectation or hope that a company will share a report of an *entire* incident, showing the context, the indicators, how they are connected, and how it was discovered, while it is still relevant. This makes an exchange actionable for a security operations team, but it also increases the risk of unwanted exposure. While companies are willing and excited to consume this incident data, they are reticent to share into these communities because they are worried about the market and reputational risk of sharing with the wrong person. Secondly, security operators need to benefit from sharing. Historically, sharing groups and structures have relied on altruism to incentivize sharing, but that is not enough. Incident exchange mechanisms must operate at the speed of business while managing risks.

EA: How important is it that sharing include an option for anonymity?

PK: Anonymity, coupled with a well-vetted community, is critical. Historically, when a CISO faces a problem, they shy away from sharing with anyone beyond their closest friend group because they are wary of exposure. The problem with such limited sharing is that the likelihood that one of their buddies is experiencing something similar and is able to offer additional insight is slim. There may be a willingness to share some data after-the-fact, but then it is too late. The company experiencing pain remains in the dark of what others may have experienced, hindering investigation and response. Others remain unaware and become prey to the same attacks. With anonymity, a company can quickly begin collaborating with a much broader set of known good guys, which may include the competition, investors, or suppliers, all without fear of being connected to the incident. Anonymity enables incidents to flow far earlier in the incident response cycle and protects against the market or reputational risk of sharing. However, anonymity is not a panacea. It addresses risk of exposure, but companies participating in an exchange must also receive an immediate benefit, such as correlation, from their participation.

EA: When information is collected by an organization, how easy is it for correlation to occur from different sources? And can this be done centrally by a third-party?

PK: The problem to-date has been that we are trying to correlate sensor-based and open-source data streams with each other and gain actionable insight. While correlating these is easy, and is done today, the resulting analysis is not actionable; it lacks context or explanation of what is happening and what steps to take. When you start with real incident data or an anomalous event, it is easier to bring in context by correlating with similar events at other companies, with threat service feeds (like VirusTotal or Farsight), and open source reports. Without a doubt this can be done by a third party, and I would argue that it *needs* to be done by a third party to be effective. Otherwise, you are only working with one part of the puzzle.

EA: Collaboration between security operations teams has been pretty thin the past few years. Do you think the necessary incentives can be provided to improve this situation?

PK: SOC teams are far more likely to collaborate if they also get something out it. Since the likelihood of identifying correlations with a particular incident increases as the contribution-base increases, I see this boiling down to a question of how to establish comfort with sharing at a large scale. At TruSTAR, we believe that protections such as anonymity, vetting, and redaction are enough to establish that comfort, and thus build out an active-enough exchange that brings value to those that contribute to an exchange. Equally important is the value back to the operators in that they need to see correlated results immediately and are notified of relevant changes or new developments immediately. With these capabilities in place, operators need a means to collaborate too. Being able to engage others over

common data sets is important. These conversations, which enrich reporting with additional context and expertise, should persist allowing operators to later return to a conversation and build narratives around events. Alerting allows operators to continue to engage with each other as conversations develop without having to have eyes on the screen all the time.

EA: What is the current state of automated ingest of threat information? Are there good standards in the industry?

PK: There are a lot of platforms that provide sensor-based data, however, these often lack the broader context needed for security operations teams to understand and act upon the information. We need to move towards incident data as a primary source and there is not a gold standard yet for automated ingest of *incidents*. But, there are supporting mechanisms in place, like STIX and TAXII that will help guide this effort going forward.

EA: Do you think global cyber security coordination and cooperation is going to be possible? It seems inconceivable to imagine America, China, and Russia sharing threat information in a friendly, cooperative manner.

PK: We live in a world with multinational companies that support customers around the globe. If we could coordinate more effectively with each other, the vast majority of attacks would be reduced significantly. But, we need to be realistic. Ensuring that you are coordinating and collaborating with trusted and vetted partners is paramount. An exchange mechanism must have the means of continuously vetting companies and also be able to discharge those parties that are not trustworthy.

Distributed, Application-Aware Micro-Segment Security

Implementing a Distributed Security System for the Virtual and Cloud Enterprise

Tim Eades, CEO of vArmour

Computer scientists have known for many years that distributed systems are more resilient than centralized ones. So with modern enterprises transitioning from centralized perimeter-protected gateways to more distributed protection, expectations are high that cyber security will improve. The resulting cloud workload protections become like mini-perimeters and this distributed, virtual approach increases the ease with which the security tools can *understand* the applications they are designed to protect.

EA: Why has it been so hard to secure data centers?

TE: Critical applications and data are protected by old architectures based on the premise of locking down these assets and limiting access. With data center transformation, adopting technologies such as virtualization, cloud, and containers, all of a company's assets have been liberated. That is, they now move freely based on efficiency, they can be transient, and they interact with a whole bunch of other systems. Legacy hardware-based security architectures were not made for this new world.

EA: What specifically is a micro-segmentation and how should a CISO use the technique to protect data?

TE: Segmentation has been around for a decade, evolving from firewalls to NGFWs. However, security requirements to support new IT models have changed so dramatically that the old ways of segmentation cannot effectively scale or secure critical assets and infrastructure. At vArmour, we patented our approach that moves controls that were traditionally at the perimeter and instead places them around

every workload or application in your environment, called micro-segmentation, which you can think of as micro-perimeters. CISOs can now keep assets with differing levels of security requirements across common, shared infrastructure, with no more silos because of security. For example, a large retailer is able to place PCI and non-PCI assets on the same hypervisor while keeping assets logically separated.

EA: Can CISO teams easily migrate from their existing security architecture to a distributed security system?

TE: Absolutely. Distributed security systems (DSSs) aren't like traditional appliance-based ones that are costly and complex to manage. Distributed security architectures instead are all software, making it easy to install and deploy virtual perimeters across hundreds of thousands of different virtual machines in times measured in hours rather than days.

EA: Are they any advantages over a distributed security system (DSS) over a traditional hardware-centric architecture approach?

TE: Everything in the data center is going to software – this includes software defined compute, storage, networking – and security is no different. The same benefits of greater agility, cost efficiency, and scale apply to a software-based distributed security system. Our customers tell us that what makes a DSS unique is the value of one integrated system: namely, no more single instance systems to manage, but rather one set of global policies across the infrastructure; no more painful hardware refresh cycles; and no more buying multiple products to get better context and correlation that you get from an intelligent system.

EA: How hard has it been for your customers to convince auditors and regulators that this new virtual method of protection is acceptable?

TE: It's been a hard road for customers to navigate with auditors. When virtualization was first introduced, it was impossible for auditors to put their arms around it. It wasn't a physical machine they were auditing since several VMs could sit on one machine. What was in scope, they would ask, the machine or the OS or both? Most auditors will want to include everything as in scope for a virtualized environment because it's difficult to have borders, segmentation, or boundaries with virtualized environments. We've partnered with our customers to help them plan ahead of new regulations since our distributed security approach can satisfy both the auditor and the customer's business requirements without the traditional friction.

EA: What other ways can a distributed security system (DSS) help CISOs with their defense-in-depth strategies?

TE: The core value of a DSS is that it's one intelligent, connected system. Because its software, additional security functions or capabilities can easily be added to the system that you would normally find in point security products. We first designed the system to provide security monitoring and segmentation capabilities. Now, we have introduced cyber deception that provides a proactive defense against attackers in an approach that removes many barriers for adoption. It is simple to deploy and cost-effective without additional infrastructure resources and secure.

Securing the Keys that Secure the Internet

Proper Protection of Keys and Certificates is an Essential Element of Global Cyber Security

Jeff Hudson, CEO of Venafi

Most cyber security professionals take for granted the essential role that digital keys and certificates play in the underlying model supporting the Internet, eCommerce, and most forms of on-line interaction at the enterprise level. However, there is a growing acknowledgement by many CISOs that this is a foundational component of cyber security protection, one of the few techniques for truly reducing their attack surface. The underlying public key infrastructure, along with world-class means for protecting keys and certificates must be viewed as among the most critically important elements of the modern enterprise security program.

EA: Do most security professionals understand the notion of trust in their cyber protection programs? It's an intangible concept, but seems so essential to a good approach.

JH: I have met with hundreds of CISOs and CIOs over the last three years, and for the most part, trust is not defined in their organizations – and consequently most organizations do not take sufficient action to protect trust. Part of the problem is that there is little understanding of what *can* be trusted, and *why* it can be trusted. The relatively simple concept of trusting people has been operationally defined within organizations by the use of usernames, passwords, or biometrics to identify individuals and trust them based on rules or policies. When I ask a CISO if a given device can be trusted, or if some application can be trusted, and how would they know, the answers I get generally indicate a lack of understanding of what constitutes trust. There is a broad based lack of understanding that keys and certificates are the foundation of trust in community.

EA: How does a security team make decisions about which certificates should be trusted?

JH: In general, they don't make the decisions. At Venafi, we support approximately 300 of the world's largest corporations, and we can say with certainty that a very high percentage do not understand how many keys and certificates are active in the network, how they got there, who owns them, and whether or not they can be trusted. This demonstrates the lack of an enterprise wide awareness and visibility system. Additionally, it proves the lack of an automated, secure lifecycle.

EA: Do you see modern cyber security teams taking PKI more seriously? It got a bad name years ago due to complexity. What types of advances have been made since then?

JH: First of all, everyone knows that certificates are everywhere, and that their use is growing rapidly. Certificates are the only viable mechanism to determine if devices or software entities should establish mutual trust, and whether they can communicate in private using encryption of the data in motion. The problem is that certificates are blindly trusted today. There are over twenty million pieces of malware, for example, that are signed by stolen or forged certificates. The bad guys know that if they sign the code, which involves marking it as authentic using a certificate, no one will check to see if the certificate is lost or stolen. In other words, they are blindly trusted. That is, if you have one, then it is trusted. And with certificate usage skyrocketing, the risks are getting worse.

EA: How do mobility and cloud services affect the overall cryptographic architecture of a typical organization?

JH: The boundary of the organization has disappeared. It used to be that everything was behind the firewall, but mobile, cloud, Internet, and IoT have completely changed all of that. There is simply no boundary. If you have a fence, for example, you can assume most things inside the fence belong. If you can't put a fence around things, how do you know what belongs? The only way is by tagging everything that belongs. That is what keys and certificates do – they identify devices and software. So the change is in the very geographically distributed nature of keys and certificates and the fact that on the Web there are many keys and certs that may belong to you and many are fraudulently belonging. The task is to have visibility throughout the Web, and inside and outside your firewall.

EA: I'm sure you've been asked this a thousand times, but how to you feel about the big debate between Apple and the FBI? Any thoughts you can share?

JH: The iPhone is very good at defending itself. The FBI asked Apple to create a version of the software that wasn't so good at defending itself, so the FBI could get into the iPhone in question. The real question is that with all of the FBI resources and capability, why didn't the FBI write the software? The funny thing is that most

people cannot answer this question. The answer is that the iPhone is designed to only run software that is signed by the code-signing certificate that is closely guarded and controlled by Apple. So the single most important piece to unlocking the iPhone is a certificate. The lesson for everyone is that certificates may be the single most important security mechanism in the world today.

Virtual Platform for Advanced Cyber Protection

Providing an Underlying Virtual Infrastructure for Hosting and Implementing Cyber Security

Alex Tosheff, CISO of VMware

Computer servers have evolved from the single CPU-per-OS implementation that has existed for so many years to a new model based on multiple virtual machines running over a single physical host. As one would expect, this model, which includes a unique piece of orchestration software called a hypervisor, creates both challenges and opportunities in cyber security. For the vast majority of enterprise data centers and networks, however, the opportunities of this shift to virtualization will far outweigh any difficulties.

EA: Alex, help us understand the recent evolution of the modern data center toward virtualization.

AT: In a traditional data center, the application, database, and storage tiers reside as separate entities, distributed on the network. This creates a significant benefit in terms of scalability, reliability, and agility in deployment and management. As these architectures grew in complexity, however, network teams have been forced to flatten the network topology – essentially allowing more and more hosts on the same network segment. The downside of this approach is that these environments become more vulnerable to attacks from the inside, due to the open nature of the architecture. It also becomes simpler for a network-based attack to traverse the network laterally (so-called East-West traffic). One solution is to deploy hundreds if not thousands of firewall rules to restrict access between tiers and servers, but this very quickly becomes overly complex to manage and creates what are called compound-policy errors in the firewalls. However, by fully virtualizing all host, server and network functions using SDN, the opportunities to keep the benefit of multi-tier architectures, while also creating significantly more secure deployments

are clear. Virtualization creates a wide variety of benefits, ranging from support for micro-segmentation, guest-introspection, services and so on. Future cyber security solutions will almost certainly depend on a virtual platform for optimal protection.

EA: Are there still data centers being built that do not employ some form of virtualization?

AT: I guess there must be some conventional hardware-oriented data centers still being built, but the vast majority of data center designers and managers are either already in a virtual model or are working toward some form of virtualization and SDN. It makes perfect sense for this, since the technology in support of virtual computing has gotten so much better and the cyber security benefits to the business, which is what you and I care about most, are dramatically improved.

EA: Is the primary benefit of virtualization reduced cost? Or do you see enhanced provisioning and improved functionality as being more substantive drivers?

AT: It's all of the above. The economics of virtualization are well understood. But clearly, if the lower costs of virtual data centers came with losses in provisioning and functional capabilities, then there might be some hesitation. But exactly the opposite is true. Data centers running virtual operating systems such as from VMware will operate more efficiently, more cost effectively, and with greater functionality such as simplified provisioning across virtual machines.

EA: Should cyber security teams worry about attacks on hypervisors? Could malware, for example, cascade across a group of hypervisors and wipe out a data center?

AT: No one in cyber security ever says that an attack is not possible. But the likelihood of this is greatly reduced by controls in the virtual operating system, by security capabilities and design approaches in the hypervisor, and by additional security controls that can be deployed for greater visibility and protection across the deployed infrastructure.

EA: Tell me how a cloud workload might be dynamically protected in a virtual operating system? Are the protections created in the same manner as the functionality being secured?

AT: We've entered an exciting time when it comes to securing enterprise workloads in private, hybrid, or even public clouds. Because we have abstracted the network, host, and security functions as software defined, we can now do some really amazing things. A key benefit is micro-segmentation – which involves creating networks with security attributes to protect at the application level. This greatly reduces the risk of attackers moving laterally through the datacenter, because every application, database, or storage server can essentially live in it's own segment, complete with network-based and host-based security. We've also taken this one-

step further. Our software-defined datacenter capability (SDDC) allows enterprise workloads to be completely portable and live in your private cloud, in public clouds, or even distributed between multiple clouds. Any time these workloads are moved, their security capabilities move right with them. A key security tenet is to make it expensive for the attacker to succeed, and this architecture does just that through high levels of distribution, management, and scale. Another very powerful capability enabled by virtualization or SDN is the ability for the network hypervisor to natively provide guest-introspection and data security capabilities. Quite simply this means with a tool like VMware NSX service composer, that you can setup a policy that provides DLP functionality at the host-level with just a few clicks. You can also define a manifest of third party integrations like anti-malware and other endpoint tools. There are so many things you can do once you abstract hosts and network functions into software-defined. It's pretty amazing if you think about it!

EA: Since virtual operating systems make services accessible to running programs through application programming interfaces or APIs, as most people know them – do you see the virtual API as the new security gateway for enterprise?

AT: They are already the new security gateway for the enterprise. This is a big deal, because it changes the way you create a cyber security architecture for an enterprise or data center. Instead of a bunch of equipment looking for packets or network activity, the new concept is to embed protections for real time application security into the cloud operating system. The result is a virtual cyber security architecture, and our team at VMware considers this evolution to be a great opportunity to provide a new virtual platform for security services. Data center managers are doing this now, and the Internet service providers are doing it with SDN deployments.