

Escaping Security's Dark Ages

Amit Yoran

[RSA Conference 2015](#), April 21, 2015

Since the beginning of time, humanity has been afraid of the dark. And with good reason. We fear the dark because evolution has hard-wired us to be suspicious of it, or more specifically, suspicious of the potential threats that may await us in the darkness. We can hear noises and see shadows, but without being able to see our surroundings, we don't know if those sounds and shadows represent danger or not, let alone how to respond.

You know what's really scary? Me. On stage with a mic. Right now our lawyers are absolutely terrified so don't be surprised if the lights stay on but my mic "mysteriously" dies.

Seriously though, my stumbling around in the dark is a pretty good metaphor for everyone who's trying to defend a digital infrastructure today. Every alert that pops up is a bump in the night, but we're often in the dark without enough context to know which noises really matter, and which can be ignored.

Before we get to the security implications of our current situation, let's talk about this world we're bumping around in. It's not that much of an exaggeration to say that we sit at a critically important inflection point not only in our industry, but in human history.

Now before you accuse me of hyperbole or exaggeration (something which I would say that our industry is all too often guilty of), consider this. The computer technology we're charged to protect has accelerated our society to heights we never could have imagined even a few decades ago. The information age has been heralded by systems with incredible computational capability, data stores vaster than human comprehension, and a speed of communications that boggles the mind. They push the very boundaries of our imagination, daring us to dream what the art of the possible might become. And they're advancing at breathless speed.

The speed of the world's fastest \$55M supercomputer in 1996 had been nearly doubled ten years later by the \$500 PlayStation 3. In 2011, a computer crushed arguably the best human contestants ever on Jeopardy!, a game requiring incredible speed of intuition and deduction, leaving contestant Ken Jennings to quip, in his final response – "I for one welcome our new computer overlords."

Even things that just a few years ago we thought of as innately human tasks have now been mastered or are on the way to mastery by computers – creative writing, emotional expression, software development, creating art, and driving cars. And even in these seemingly uniquely human tasks, computers are proving themselves better than us! In the first million miles driven by Google cars, the only two accidents that occurred were caused by human override. We are

taking a backseat to our own technology, and apparently we don't drive any better from that position either.

One doesn't have to imagine very far into the future to see the logical outcomes of genetic research, microscopic robotics, self-organizing networks, self-programmables, sensory enhancing technologies, and 3-D printing – of even organic material. While we're not there yet, it is no longer some flight of science fiction fantasy. In fact, it's extremely likely that in the next few years technology will be capable of accelerating its own development, and in short order creating organisms we would be hard pressed to call anything other than life itself. Without a doubt, we are at an inflection point for humanity, where technology will control its own destiny, the results of which we cannot predict.

We stand at the doorway of The Age of Technology Enlightenment, where we have already become completely reliant on computers in every aspect of our personal and professional lives.

Unfortunately, as we look through the doorway amazed at what lies ahead for our civilization, we still surely stand in the Dark Ages of Information Security.

2014; yet another “year of the breach”, or have we agreed to call it the “year of the mega-breach?” That might at least connote that things are getting worse, not better. And you don't have to be much of a visionary

to realize that 2015 will be called the “year of the super-mega breach.”
At this pace we will soon run out of adjectives!

The largest enterprises with the most sophisticated, “next-generation” security tools were not able to stop miscreants from making off with millions of dollars, personal information, and sensitive secrets and damaging reputations. And damn it, they even messed with Seth Rogen! That’s just not cool, man! By the way, we invited Seth to come speak at the conference, but he declined. Next year we’re going to try getting this Supreme Leader guy to come.

2014 was yet another reminder that we are losing this contest. The adversaries are out-maneuvering the industry, out-gunning the industry, and winning by every measure.

The general purpose computing theory that we rely upon is fundamentally broken. Alan Turing was wrong. Modern computers and interconnected networks are not deterministic machines or finite state automata at all. Similar to humans, an infinite and unpredictable number of inputs and influences drive them to a seemingly infinite number of follow on states. Those of you who are married know exactly what I’m talking about – you can never predict how your spouse might respond to input.

Simply put, and for all practical purposes, we can neither secure nor trust the pervasive, complex, and diverse end-point participants in any

large and distributed, computing environment; nor the combinations of protocols and transports through which they interact. That is the situation we are in today.

If you have been in security for a while, you probably agree that the industry has promoted a defensive strategy that aligns with a Dark Ages mindset – to keep the barbarians away, we’re simply building taller castle walls and digging deeper moats. Now, I suppose in fairness, we have built some *next generation* walls with application-aware passageways through them. But as anyone who looks at our reality can tell, taller walls won’t solve our problem.

Put another way, it’s like we’re working from a map of a world that no longer exists; and possibly never did.

As some of you know, I’m a West Point graduate (yeah, my mom has a hard time believing that too). I remember one time during my land navigation course, I was completely lost. No matter how I compared the map to my surroundings, I couldn’t make heads or tails of my location. I turned to my instructor, a crotchety old Special Forces sergeant, and explained that the map didn’t line up with the terrain; there was supposed to be a mountain over there. His first response, if I remember correctly, was something like, “Cadet, you couldn’t find your ass with a map and both hands.” And if you know me, it took every ounce of willpower not to smile, grab my buttocks, and say “found it, Sergeant!” And then he said either the terrain was wrong or the map was wrong . . .

and I could tell by the way he said it, that he didn't think the terrain was wrong.

It's clear in security that we haven't been able to find what we're looking for even with a map and two hands. The map we're looking at doesn't fit the terrain but we keep pretending it does. Perhaps we're just hoping that *this time* the terrain will magically change. Perhaps *this time* perimeters will protect us. And before everyone starts shouting, "we know perimeters aren't sufficient," the perimeter mindset is still with us; we're still clinging to our old maps, hoping the terrain is wrong.

Beyond the irrational obsession with perimeters, the security profession follows an equally absurd path to detect advanced threats. Monitoring is performed with signature-based intrusion detection systems and anti-malware products. It's not that perimeter and preventative measures are bad in and of themselves, it's that they are limited by experience. They have to have seen a threat before or have been taught about it to be able to detect it. We've all heard that the threats that matter most are the ones you haven't seen before. These tools by definition are incapable of detecting the threats that matter to us most.

Nonetheless, many security professionals base their programs on the futile aggregation of telemetry from these virtually blind IDSes, AV platforms, and firewall logs, implementing the glorious and increasingly useless money-pit, known as the SIEM. I know it didn't surprise many of

you when last year's Verizon Data Breach Investigations Report asserted that less than one percent of successful advanced threat attacks were spotted by SIEM systems. Less than 1%. The terrain has changed but we're still clinging to our old maps. It's time to realize that things are different.

Isaac Asimov once said the most exciting phrase to hear in science, the one that heralds the most exciting discoveries, is not the discovery itself – “Eureka!” (I found it) – but rather the awareness of a problem – “Hmmm, that's funny.”

I'm sure the scientists who discovered Viagra's benefits probably uttered those exact words. They had been testing the active drug in Viagra in the hopes that it would become a block-buster cardiovascular drug that lowered blood pressure. During the clinical trials, they found that people didn't want to give back the medication because of the, umm, interesting side effects. “Hmm, that's funny.”

In our far less exciting case, “Hmm, that's funny.” The strategies and systems upon which the security profession has been relying don't produce the result we expect. It is time for a renewed sense of exploration, awareness, and understanding. It's time for security to escape the Dark Ages and pursue our own Age of Enlightenment.

Without foundational shifts in computer science, fundamental research into different building blocks, and different connective tissue – which

isn't happening anytime soon – we're going to be dealing with these challenges for a long time to come. The sophisticated barbarians are *already* inside the gates. Not only are they inside the gates, but they've raided the liquor cabinet, and are walking out with anything that hasn't been bolted down.

Even organizations that invest massive amounts of money in protecting themselves are continuously being compromised. Our Asimov moment should be the catalyst to change our industry's mindset and to start thinking and doing things differently.

No doubt some of you are saying, "Intruders are on the inside? Tell us something we don't know." The reality is that whether we claim to know it or not, our actions don't reflect it. It's like the scene in the Matrix when Morpheus says to Neo there's a difference between knowing the path and walking the path. We say that we know the perimeter is dead and that the adversary is already inside the gates, but we aren't changing how we operate.

So how do we re-program ourselves for success? What is the path forward? Let me share five thoughts on navigating our new terrain, based on my conversations with many of you and my own experiences over the past 20 years doing advanced threat detection and response work – from my time in the DoD in the 1990s through to today including the observations of RSA's incident response team.

First.

Let's stop believing that even advanced protections work. They do, but surely they fail too. Here's the news flash that has underwritten each and every spectacular intrusion we read about on a daily basis and countless others that remain undiscovered and unreported – and that's that a well-resourced, creative, and focused adversary is going to get into your environment. Every modern nation-state and every organized criminal enterprise operates aggressive intelligence collection or monetization schemes online. They enjoy limitless bounty with near perfect impunity.

You'll see many promises made this week – expect to see more big data solutions, solutions to IoT, “fire and forget” analytics, and all sorts of other buzzwords, but challenge yourself and challenge us vendors - does this really help or is it yet another wall that will inevitably be breached? We're seeing analytics-resistant malware that can evade detection by sandboxes and other advanced systems. No matter how high or smart the walls, focused adversaries will find ways over, under, around, and through.

Second.

We must adopt a deep and pervasive level of true visibility *everywhere* – from the endpoint to the network to the cloud – if we have any hope of being able to see the advanced threats that are increasingly today's

norm. Consider the Stuxnet, Equation Group, and Carbanak intrusion sets and countless other sophisticated campaigns. One of the defining characteristics across all of them is their stealthy nature. Until written about, they were virtually undetectable because they bypassed traditional defenses. Even now many organizations operate completely blind as to whether they are victim to these published techniques.

We need pervasive and true visibility into our enterprise environments. In reality, I'm describing now what SIEM was meant to be, or rather what it should be. You simply can't do security today without the visibility of both continuous full packet capture *and* endpoint compromise assessment visibility.

Within our digital environments, we need to know which systems are communicating with which, why, any related communications, their length, frequency & volume, and ultimately the content itself to determine what exactly is happening. These aren't nice to haves, they are fundamental core requirements of any modern security program. If you don't have that level of visibility and agility in place, you're only pretending to do security.

Traditional forms of visibility are one-dimensional, yielding dangerously incomplete snapshots of an incident let alone any semblance of understanding of an attack campaign. Without the ability to rapidly knit together multiple perspectives on an attack, you'll never fully understand the scope of the overall campaign you're facing.

Frequently, sophisticated adversaries are executing attacks using multiple tactics in concert, often even from separate attack groups to assure persistent access. The single most common and most catastrophic mistake made by security teams today is under scoping an incident and rushing to clean up compromised systems before understanding the broader campaign. In fact, let me say that again. The single greatest mistake made by security teams today is under scoping incidents and rushing to clean up compromised systems before understanding the true scope of compromise and possibly broader campaign. Without fully understanding the attack, you're not only failing to get the adversary out of your networks, you're teaching them which attacks you are aware of and which ones they need to use to bypass your monitoring efforts. As Sun Tzu said, "If you know the enemy, blah, blah, blah." Sorry, it's not a security conference until someone quotes Sun Tzu so I figured I would get it out of the way. Forget Sun Tzu. If you want to know about combat, spend a day with my kids. Teaching your adversaries how to better attack you is what my kids would call an "Epic Fail!"

And I'm not just standing up here and saying "buy RSA gear." I'm the first to admit that we need to go further than what is available today. We're on a journey to full visibility. Our environments, business practices, and adversaries continue to evolve and so must we.

Third.

In a world with no perimeter and with fewer security anchor points, identity and authentication matter more than ever. Consider the evolution of today's threat actors and their tools, tactics, and procedures (or TTPs – I got that in for those of you playing buzzword bingo). Today's anti-malware solutions are great. Buy them and use them. But don't mistake an anti-malware solution for an advanced threat strategy. The Verizon Data Breach Investigations Report shows that malware was the primary attack vector in less than half of the advanced threat breaches. In cases where confidential data was disclosed, the most popular method used was Web application attacks. And in those cases, 95% of the time, attackers used stolen credentials and simply walked right in. The Verizon report talks about how often users' mistakes, not sophisticated malware or hacks, open the gates to the adversaries. It reminds me of a post I saw once that said, "Who needs zero days when you've got stupid?" Of course, "stupid" is a bit harsh since our users are only human and the right social engineering can get even the most knowledgeable to click on that link or cough up their credentials.

At some point in the campaign, the abuse of identity is a stepping stone the attackers use to impose their will. The creation of sysadmin or machine accounts or the abuse of over-privileged and dormant accounts facilitates lateral movement and access to targeted systems and information (lateral movement – ding – buzzword). Strong authentication, and analyzing who is accessing what, can identify attack campaigns earlier in the kill chain and make the difference between successful response and unmitigated disaster. Don't make the mistake

of just trusting the actions of the trusted; those are the very accounts and users most targeted and of which we should be the most suspicious.

Fourth.

External threat intelligence. This is a core requirement as well. There are incredible sources for the right threat intelligence for your purposes from vendors like CrowdStrike, iSIGHT Partners, ThreatGRID, and others. And beyond private vendors, there are organizations like ISACs. Threat intelligence should be machine-readable and automated for increased speed and leverage. It should be operationalized into your security program and tailored to your organization's assets and interests so that analysts can quickly address the threats that pose the most risk. And for God's sake, do away with pdf and email sharing and response coordination. In fact, we've seen adversaries compromise mail servers specifically to monitor sysadmin and network defender communications. Ouch.

And finally.

You must understand what matters to your business and what is mission critical. This asset categorization isn't the sexy part of security but it is critical to helping you prioritize the deployment of limited security resources for the greatest possible impact. You have to focus on the important accounts, roles, data, systems, apps, devices – and defend what's important and defend it with everything you have.

These ideas can work. They do work. We've seen the difference it makes when organizations take these approaches to security. We see customers understand the attack campaigns that have been running in their environment for months or longer - often right under the noses of their protective measures. In one incident response effort, we discovered breach artifacts that were in place for seven years. Seven years. With these ideas and agile mindsets, our teams are even catching attackers red-handed, and disrupting their ability to exfiltrate data and achieve their goals.

I'm not saying we have all of the answers - far from it - there are resource challenges, there are skills challenges, there are legal challenges. But we are on a path to changing a paradigm under which our industry has operated for decades. And at RSA, we're starting with ourselves. We're re-engineering RSA across the board to enable us to deliver on this vision. This time next year, we won't be the same RSA you have known for decades. As an industry, we are on a journey that will continue to evolve in the years to come through the efforts of all of us. I'm reminded of an old story from the age of exploration when many maps weren't yet complete. As the story goes, a captain was sailing his ship and reached the edge of his map. He sent word back to his commanders, "Have sailed off the map. Am awaiting instruction."

We have sailed off the map, my friends. Sitting here and awaiting instructions? Not an option! And neither is what we've been doing -

continuing to sail on with our existing maps even though the world has changed.

What I'm describing is not a technology problem. We have the technology today to provide true visibility. Strong authentication and identity management solutions are readily available. We have great threat intelligence and insight into sophisticated adversaries. And we have systems that map and manage our digital and business risk.

This is not a technology problem. This is a mindset problem.

The world has changed and trust me, it's not the terrain that's wrong.

Thank you.