SANS

WHAT WORKS™

WhatWorks in SOC/NOC Integration: Improving Time to Detect, Respond and Contain with ExtraHop Reveal(x)

ExtraHop

Rise Above the Noise.

WhatWorks is a user-to-user program in which security managers who have implemented effective Internet security technologies tell why they deployed it, how it works, how it improves security, what problems they faced and what lessons they learned. Got a story of your own? A product you'd like to know about? Let us know.

www.sans.org/whatworks

The 2019 SANS Security Operations Center survey showed that SOC managers list lack of integration between security and IT and network operations as one of the major obstacles to both their effectiveness in dealing with evolving threats and their ability to efficiently protect the business in constrained budget environments. Less than 40% of SOC managers say that the SOC and the NOC are effectively integrated, and those who report higher levels of integration show improvement in reducing time to detect, respond and contain.

During this SANS What Works webcast Mitch Roberson, Director for Enterprise Systems at Curo Financial, will provide details of the selection and deployment of ExtraHop's Reveal(x) to increase visibility into network traffic, gaining detailed and timely insight into performance and security issues and crossing organizational siloes by using a common tool and dashboard for application owners, network administrators and security analysts.

Join SANS Director of Emerging Security Trends John Pescatore and Mitch Roberson to hear details on the selection, deployment and experience using ExtraHop. The webcast contains a discussion of lessons learned and best practices as well as detail the metrics used to demonstrate the value of improved email authentication and trust.

ABOUT THE USER

Mitch Roberson, Director of Enterprise Systems for Curo Financial, is responsible for the majority of the company's applications, servers, storage and hardware. He serves on the incident response team at Curo and works closely with the security and the networking teams. Having worked as a consultant at multiple VAR's as well as Microsoft, Mitch has seen a multitude of environments and worked with network, systems and security teams. This has allowed him to broaden his knowledge in many areas of IT. Because of this broad experience, Mitch has an almost fanatical desire to have visibility into his environments. His passion is to learn how applications communicate so he can decrease mean time to resolution and uncover malicious activity as early as possible.

ABOUT THE INTERVIEWER

John Pescatore joined SANS as director of emerging security trends in January 2013 after more than 13 years as lead security analyst for Gartner, running consulting groups at Trusted Information Systems and Entrust, 11 years with GTE, and service with both the National Security Agency, where he designed secure voice systems, and the U.S. Secret Service, where he developed secure communications and surveillance systems and "the occasional ballistic armor installation." John has testified before Congress about cybersecurity, was named one of the 15 most-influential people in security in 2008 and is an NSA-certified cryptologic engineer.

Q Mitch, tell us a little bit about yourself, the role you play at Curo and maybe a little bit about Curo.

A My name is Mitch Roberson. I'm the director of enterprise systems for Curo Financial. I'm responsible for the majority of the applications, the servers, the storage, the hardware. I also am on our incident response team and also heavily involved with the security team as well as the networking team. I report to the VP of global infrastructure, who reports to our CIO.

Q Is there a chief information security officer? Where do they report?

A We actually have a director of information security, and he reports to the VP of global infrastructure as well, which is also my boss.

Q Tell us a little bit about Curo Financial.

A Curo is a company that runs locations in Canada and the U.S. We have about 400 retail locations, and we have a very large web presence in both Canada and the U.S. We serve underbanked customers with loans and financial services. We have three data centers, one in Canada, two in the U.S.

ExtraHop was the only one that said they didn't have a problem with any of the ten use cases. So we brought them in to do a proof of concept.

Q This What Works will be about your selection and use of ExtraHop's Reveal(x) product. What was sort of the business or performance issues? What started you in looking at products like this?

A Curo is a company that grew from a very small organization to an enterprise company rather quickly by most standards. When I got here, there were a lot of things that people couldn't answer for me about how things worked. So I started pushing for visibility into the environment pretty early.

We brought Riverbed in, and I worked for three years with their product. But, I always felt like I was missing details out of the network side of things that I thought we needed and should be able to get.

When Riverbed came up for renewal, I decided that I would go out and look at multiple vendors. I picked several, including Riverbed as the incumbent. ExtraHop came in, NetScout came in and there were a few others that I contacted and sent a list of ten items that I wanted to see out of their product. I had a pretty good set of use case scenarios that I built out and I wanted to see, at a minimum, those ten things. If they couldn't do it, I pretty much told them that I didn't even want to bring them in for a proof of concept. ExtraHop was the only one that said they didn't have a problem with any of the ten use cases. So we brought them in to do a proof of concept.

Q Did your use cases include security use cases, or did you bring people from the security side? Quite often we see the network operations looking at tools and security operations looking at tools and not necessarily always doing it together. How did you deal with that?

A We're a little unique because we interact so tightly together with everybody and I have a background in dealing with security issues.

When we started bringing this in, part of it was for that security visibility and knowledge. I was really concerned about things like exfiltration and ransomware attacks. We had recently suffered a very small ransomware attack that wasn't very successful, which is good. I wanted to get more visibility into that as well as many other things. Out of ten questions we asked of the vendors, around seven of them were specifically around security things that we wanted to see.

Q Were you evaluating them solely on paper? Did you do bake-offs, run competitors in parallel, etc.?

A We ultimately never got to the point of bringing them in to do a bake-off because with the ten use cases that I had, very few of the vendors would actually commit to all of the use cases. ExtraHop was the only company that responded, "Yes, we can do every one of those use cases that you have with no problem." We brought Reveal(x) in for a POC, I installed it and they delivered on nine out of ten in the first 48 hours. It took them another two days to deliver on the last item. I added about 20 more use cases after that, and they still hit every single one of them.

Q What is involved in installation? Are you tying the switch SPAN ports, installing taps? Walk us through what it took to get installed and get started.

A We already had a small tap infrastructure because of having worked with Riverbed in the past. I had already deployed GigaVUE from Gigamon for east-west traffic and had physical taps in place for north-south traffic.

All of this was being sent to a Gigamon packet broker so it was easy to just take a feed and send it to the new Reveal(x) platform.

Once we saw the value of the data, we did purchase additional taps and they really didn't take a lot of effort other than getting the networking guys to suffer the short outage. The actual hardware and software deployment for Reveal(x) was extremely easy. As soon as we started sending it data from the taps, it started ingesting and providing usable data within less than five minutes.

Q What does your architecture look like? You said there are three data centers. Is there an appliance at each data center, then some sort of consolidator?

A We have a Reveal(x) appliance at each data center. Our Canada data center doesn't have Reveal(x) in it yet. It has the performance appliance. We have a packet broker that receives packets from the physical taps or spans or GigaVUE appliances. At each respective data center the packet broker consolidates and, in some cases, duplicates the packets. Then it forwards those packets to the Reveal(x) or performance appliance in the respective data center. We use ExtraHop's ECA appliance that sees across the Reveal(x) platforms in both data centers. Basically, we get an aggregated view. For most end users, they don't know the difference. They don't care which data center it's in or whatever else is there. They get one view, one place to go to, and they see all that information rather easily.

Q What are those views they're getting? Is it purely raw data? Are there patterns, alerts, alarms, threshold? How are you using that data?

A We have everything from alerts to graphs of all sorts. Basically, a lot of metrics. The way we're designed is we've got all the metrics that are ingested and pulled through by the Reveal(x) appliance, and there are detections that are coming out that are alerting us. We have some of the more critical detections going to our SIEM and an in-house–built Hive case management solution.

Our security team gets the notification that a detection has occurred, and it goes to the Hive. That Hive immediately goes out and pulls additional information from the SIEM and can pull stuff from from several different locations to aggregate the data. In the email that the user gets, they may have Reveal(x) data, they may have data from the SIEM, they may have data from other tools, etc. They may have data from several locations giving them information about the two IPs that were involved in that detection or the multiple IPs that were involved in that detection. So we can quickly decide, "Do we need to continue to work this as a case, or do we need to do something else with it?"

Q On the application owner side, how does that work?

A All business applications have metrics associated with them that we've put into ExtraHop in Reveal(x). We get alerts on everything from DNS—I can follow DNS throughput. I can follow DNS requests versus responses. We can look at SIEM messages and application errors.

We've done some really cool things where we're alerting on errors from our web servers and getting specific notifications about when they're having problems. People are able to react efficiently to those because in the alerts, it gives us detailed information. Is it coming from one client or two clients? Is it affecting 20 clients? They can go drill in.

The actual hardware and software deployment for Reveal(x) was extremely easy.

When you see a web service have problems, you can go in and watch it on the dashboard and see that it had errors for two minutes. We get an alert on it. But if the errors have died off and we don't see it anymore, they can drill in and find out it was one or two customers or one or two clients that had that problem specifically. So they'll send out a notice saying, "Hey, this only affected two clients. We're continuing to monitor." So our NOC/SOC is monitoring those kind of problems as they come around.

Q Often, these types of alarms could be an application or network performance issue, or some form of cyberattack. How do you use ExtraHop to make that determination?

A We see that a lot more than people think. Years back the security guys would always talk about the fact that all they do is either allow good traffic or block bad traffic. But with all today's intelligent firewalls that are application aware, as well as the new advanced malware detection, all these different things that they have out there can sometimes slow down traffic or cause unique traffic problems.

We operate under the premise that we can't be siloed anymore. The security team gets detections. The application team also has access to see those same detections. So they're always paying attention to it from a security perspective. At the same time, as we're looking through it, we start to know it's a performance issue, we start to let people know rather quickly when those things happen.

A lot of times, it is performance related. It's not as often that it's denial of service or other attack. They may be self-induced denial of services by misconfigurations. It's amazing how many times we can find that and get to root cause very quickly. It's also brought all the teams together. Our networking team, our security team and our applications team are so willing to work together because of the visibility that we have and everybody's talking at the same level.

It's also brought all the teams together. Our networking team, our security team and our applications team are so willing to work together

That Layer Seven data that we're pulling off the wire provides so much valuable data to anybody and everybody out there if they actually know what to look for and spend a little time learning about it. And I don't care whether you're a member of the security team, a network engineer or a developer, the data that's there provides a tremendous amount of value.

Q How has the performance been from false positives and false negatives with those alerts?

A Because we're a development shop, our developers do a lot of really interesting things, which we want to happen—it is critical to staying competitive. So, we have a few more false positives than what I think most companies would see just simply because our developers are going out into GitHub and pulling stuff down and trying different things.

The cool thing is that the Reveal(x) tools have an excellent way of allowing us to hide that type of alert. We can still see them if we want. But we can say, "Ignore or deprecate anything coming from this box," because we know that's a development box and sometimes they're doing that to test their own application. We've gotten it down to where today, we probably have 20 to 30 alerts that we really investigate on a regular basis. And that's pretty good compared to the large number of detections that sometimes are generated but hidden from our day-to-day view. We are able to focus on the real issues and not get buried in noise.

Q You have the sort of signatures or patterns and so on that are updated by ExtraHop. Did they give you tools to where you can generate your own signature or custom alerts?

A Absolutely. Now, they're not really a signature-based solution. They're more of a behavior-based solution. So a difference in anomalies, a change. They're looking at the things that are really common, like a ping sweep. A machine has never done a set of pings before and all the sudden pings the entire subnet. Those things come out, and they jump out really fast at you from network data. Things like denial of service attacks, those things are very obvious on the wire. There are certain parameters that always go into that on the wire. So they're looking at that with behavior-based anomaly detection.

The cool thing is you can build your own. Pretty much if you can see it on the wire, you can pull it out in Wireshark, you can build your own detections and your own alerts based on that. It's pretty fascinating how well it can be done.

Q Are you doing long-term storage of packet data or just metadata?

A We're doing about three to five days of packet data right now. We'd like to extend that only for forensic purposes. We're actually moving many of our people away from packet-level data because the information that we get out of what's called records, which is the metadata, and the metrics themselves are where the real value is because we can build new custom metrics that contain the data we need.

Let's say we find an attribute that we want to map on a regular basis because we know it's a bad thing, We know it happens every so often. We can actually write a trigger to grab that data, put it in as a custom metric and save the information of who started it, who stopped it, who was involved in it, etc. And those metrics we can store even easier than metadata. So 30 or 160 days' worth of metrics is fairly easy to do in this environment.

We also have custom triggers that save metrics. That helps move our analysts away from going to packet data all the time. Not to say we've gotten rid of it. We still go into packet data. On average, I used to spend up to six hours a day in packets. For the most part, since we've gotten Reveal(x) in place, I probably spend anywhere from two to four hours a week in packet data.

Q What sort of metrics do you report to management to say, "Here's the business benefit we're getting out of this investment."

A From our perspective, the metrics are twofold. One is faster time to determine severity of the security risk. If something hits us, we can tell very quickly. We can get to root cause roughly 90 percent faster than ever before. But, the cool thing is because this tool is used across so many different environments, our IT operations team uses it, our application teams use it, our security team uses it, our network team uses it. The value shows up almost daily from almost any one of the groups.

They did a fantastic job all the way through. Their support has been great.

We can see how many detections we've had for a week. We can see how many problems we've had for a week. And we can show the value in that. We also can show something else. We've had multiple cases where we're mandated by PCI or somebody else to make a change in our environment to secure it better. In the past we would make those changes and often break things or, in some cases, we'd make the changes recommended and six months later, we'd have a penetration tester in here and find out we still had the same problem.

The value is we can see those changes as they happen and see what happens. For example, trying to get rid of TLS 1.0—everybody says do it. But I can tell you exactly what servers still are responding with TLS 1.0 today. Getting rid of weak cipher suites, I can tell you exactly which servers still have weak cipher suites on them. And so I can go after those specifically. SMB Version 1, it'd be a dream world if everybody could turn it off all the way, right? That's what everybody wants. They talk about it all the time. But I can tell you that we only have about six servers on our environment that are responding to SMB Version 1. And here in a few weeks, we'll be able to turn off globally SMB Version 1 because we've been able to surgically go after those that are still responding and figure out why and what's hitting them.

From that perspective, our management has started to really believe in the value of Reveal(x) because we're making changes that don't take down the environment, that don't cause problems. We're proving our vendors wrong on a regular basis when they say, "The problem's on your side, not ours." But, "I hate to tell you, we can show you exactly right here in these metrics. If you need to, we'll go to packets. We'll show you exactly where the problem is." When you can do that kind of thing, it becomes second nature. And your management just expects it. It's a fantastic tool that they really like.

Q Based on what you know now, are there any lessons learned, things you would have done differently that you could pass on?

A We were operational with Reveal(x) shortly after it came out. We've been using it now for one or two years. Probably the biggest thing is, you've got to get rid of the siloes. You've got to get the application guys talking with the security guys. You've got to get the networking guys talking with the security guys. The value of the tool comes when you can actually figure out your environment and build the groupings right.

We group servers based on different attributes. ExtraHop does a really good job doing this. But there's other pieces to it that become really important that are related to your company and the way the business does things. And that's when the value comes in. But you've got to have somebody that's willing to take the time to dig in and understand what this is.

People put it in, and they just start looking at the detections or looking at the few dashboards that ExtraHop provides, which are really good. I can show where we've made significant major changes to our environment. I've got over 100 different scenarios where we've made changes just based on the data that we've pulled out of Reveal(x) over the past year to 18 months.

Q **What sort of support do you use from ExtraHop, and how do you rate their support?**

A Support has absolutely been fantastic. They did identify a consultant for us that they were willing to send out. Because I had gotten it deployed so fast, we decided to use it basically as remote professional services to help with more in-depth things that I needed. They built some of the more advanced triggers that I really wanted for our environment.

They did a fantastic job all the way through. Their support has been great. Like any company, they have challenges. They do occasionally release software that creates a few problems. But the interesting thing about that is they have a fantastic support pack. I run the support pack, upload it, and 99 percent of the time, they know what the problem is without ever having to remote back in or call me again. And usually, within 24 hours, they've got a fix for me. I would say probably 95 percent of the time, they've had a fix for me within 24 hours of any issues that have stood out.

Q **Of the three data centers, do you have any applications that you're not running out infrastructure as a service, the AWSs or Azures of the world? If so, how did you pull them under the covers with this?**

A We're actually looking at moving some workloads to AWS. We have not done that yet. We backed off from doing this in the past because we just couldn't get the support to handle all our problems fast enough. We think we now have a methodology to do that. And part of the reason we're looking at doing this now is because Reveal(x) is being offered in the cloud. With AWS, they have the tap infrastructure, the virtual tap, basically virtual mirroring available now. We'll be able to put Reveal(x) in place as we move forward. That's actually an objective that we're working on this year is to get some of our stuff moved into AWS.

Q **Any final thoughts you think are important?**

A It's an interesting tool that provides a visibility that I've never been able to get to. It's visibility that I've always wanted, but I've never been able to get to. And I've worked a lot of different places and tried a lot of different things with a lot of different customers.

I'll give you a great "for instance." If you take a dashboard and you have a set of numbers on the dashboard and says that I have 10,375 errors in the last 45 minutes, does that really excite you about anything? Well, it depends. But if I can take that on the fly, change that graph to a line graph and see that that 10,885 errors happened all in a two-minute time period, that changes the dynamic. It really comes down to the way that ExtraHop has been able to adapt and display the data for us and give me the flexibility to look at it from different points and different angles is one of the big values of it. And it doesn't matter whether it's a detection, a security detection or some sort of performance detection. That is one of the biggest things that I've never been able to get out of wired data. Now it's just easy. It's just there for me.