How VCU uses FireEye for Advanced Threat Detection and Prevention

A Real Case Study with Virginia Commonwealth University

SPONSORED BY

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

ABOUT VCU

Virginia Commonwealth University is a major, urban public research university with national and international rankings in sponsored research. Located in downtown Richmond, VCU enrolls more than 31,000 students in 222 degree and certificate programs in the arts, sciences and humanities. Sixty-seven of the programs are unique in Virginia, many of them crossing the disciplines of VCU's 13 schools and one college. MCV Hospitals and the health sciences schools of Virginia Commonwealth University comprise the VCU Medical Center, one of the nation's leading academic medical centers. Visit www.vcu.edu for more information.

ABOUT THE USER

Dan Han is the Information Security Officer for Virginia Commonwealth University and is responsible for the development and management of the information security program for the University. He has close to 15 years of experience working in various roles within IT, ranging from application development to infrastructure management. He spent the majority of his career working in the higher education and healthcare sectors, and has been working in the information security field for nearly 10 years. He specializes in information security architecture and security risk and compliance management. In addition to various industry recognized IT and security certifications, Dan holds a MS and MBA in Information Systems and IT Management.

SUMMARY

A University with a centralized Internet connection but decentralized PC operations found that it experienced a high level of malware events on users' PCs. The university's security team decided to evaluate network-deployed advanced threat detection solutions that could inspect traffic at the Internet border point to address the problem. After a competitive "bake-off," they selected technology from FireEye. The FireEye Threat Prevention Platform gave VCU visibility into malware that existing anti-virus (AV) solutions were not detecting. This allowed them to respond more quickly to malware events – before major damage was incurred. The FireEye product was integrated with VCU's SIEM product for day-to-day reporting and monitoring. To deal with the high-speed (10G) network speeds at VCU, the University eventually upgraded to three FireEye appliances in a load-balancing configuration.

Q Dan, what is your position at VCU and your responsibility?

A I'm the Information Security Officer at VCU. My role is the Head of Information Security for the University. I'm in charge of the security strategy and architecture for the University. Operations also falls under my group, but we are operating in a centralized/decentralized environment. So, there are certain areas of operations that do fall into other groups.

Q What factors or problems drove you to look into a technology like FireEye.

A One of the big challenges we had several years ago was the fact that there were many malware infections out in our departments. A lot of our desktop support is decentralized, and there are various ways for each of the schools to manage their own desktops. Because of that, there was quite a bit of variance at the time in the management of desktops, and we often lacked visibility from our anti-virus software. Our anti-virus software wasn't as effective as in years past. One of the things we looked at was trying to determine the malware infections that may be affecting our environments from a network layer, and seeing if there were opportunities for us to block them. It prompted us to start looking at FireEye.

In terms of accuracy and just the pure number of infections, FireEye far-exceeded its competitor when tested.

Q So, you had a good handle on the extent of the problem and then looked for ways to address the problem?

A I believe it was a bit of both. We had a feeling that a problem was rampant within our environment, simply because we do see things on the network, such as malware callback traffic going to known malicious IP addresses. It's very hard for us to see everything; while we care about our faculty and staff machines, we also have about 30,000+ students' computers on campus, not to mention the wireless mobile devices. At times, it's difficult for us to find the source of the infection and who is infected; therefore, it is hard for us to prioritize.

We were able to better detect different virus infections

Q How did you get going and get budget approval to do something about this?

A One of the things we were able to show management was that based on our network traffic, we were looking at a large amount of callback information going back out to known malicious URLs. We were able to quantify the amount of our traffic associated with those types of activities. We said, 'Yes, this is a real risk.' Aside from that, we looked to our internal desktop support group to obtain some numbers on how many virus support tickets we were seeing per month. From the information we gathered, we were able to make a solid case on why something like a network-based detection system would benefit us.

Q Walk us through what you did from there, how you looked at solutions and evaluated them.

A Basically, our goal was to come up with a more centralized way of dealing with malware in our decentralized environments. We looked at a couple of solutions: FireEye and a direct competitor of FireEye. We did a head-to-head bake-off between the two solutions, looking at the raw number of infections detected. At that time, from a detection perspective, FireEye was a far superior product. So it was a no-brainer for us to select FireEye. In addition, our peer institutions were also evaluating FireEye. One of them had recently implemented FireEye, and was pleased with it—that contributed to our decision to select FireEye.

Q Can you provide some details on how you did the bakeoff?

A Both products were hanging off of a network tap on our network Internet border and receiving the same type of information from the tap. Based on the information coming from devices, we can determine whether there are any malware infections or, malware downloads happening. In terms of accuracy and just the pure number of infections, FireEye far-exceeded its competitor when tested.

Q The devices were looking at HTTP and everything else on the network?

A Yes. They were looking at HTTP, IRC, and other related types of traffic. One thing that we found quite interesting was the difference in activities across the different protocols between the two products.

Q What about false positives? Did you see different levels of false positives?

A False positive accuracy was good, but there were definitely more false negatives from FireEye's competitor than FireEye itself.

Q How did you go from the bakeoff to operational use of the FireEye product?

A Once we decided that we were going with FireEye, we kept the test appliance as our production box. Soon after FireEye was in place, we were able to integrate it with our SIEM (Security Information and Event Management) product. FireEye was then reporting its detections back to the SIEM; the SIEM itself is correlating the FireEye detection data, along with data from our anti-virus system as well as our network flow data and log data from different servers and network appliances. That definitely helped us to paint a more precise picture in terms of threats against our environment.

...we were able to show some significant improvements in proactively detecting malware infections

We were able to better detect different virus infections and different threats, and to better classify those threats based on the asset profiles themselves. With this said, it is imperative for any organization looking to implement a network-based security appliance to properly scale the implementation so that appropriate traffic can be monitored and triaged. This means identifying the appropriate network traffic to monitor, identifying the correct location to place the appliance, and ensuring that the appliance(s) are setup in a way to handle the data throughput.

Fast forward a couple of years down the road. What we found was a single FireEye appliance was not enough to handle the traffic within our environments. On any given day, we're looking at 2 to 2.5 gigs of throughputs during peak times. What we ended up doing was working with our network tap to configure a load balancing setup. We then acquired two additional FireEye appliances, with a total of three gig throughput capability. We load balanced all of the traffic across the three.

We were able to actively detect infections and take a more proactive approach to address malware infections.

Q How long did all that take?

A We first brought in FireEye in 2011. In 2012, the FireEye configuration was in full production. Since then, we have made additional improvements to the architecture to accommodate the growth in network throughput.

Q Now that you've gone to that configuration are you still doing detection only? Are you doing any TCP reset?

A We tried the TCP reset option in late 2012. We weren't so sure how reliable the TCP reset is in our environment, because from our perspective, there may be a race condition due to the complexity of our network. We are still triaging these cases and handling them with staff time. Down the road, that's where we want to go to blocking some of these active attacks.

Q Typically universities have lots of different Internet connection points. Are you centralized as far as Internet connectivity, and can you watch everything there? Or are there other Internet points?

A The Internet connection is centralized. Fortunately for us, the entire University network and networking activity is centralized under one umbrella. It was actually much easier for us to deploy this at a central level, and that's why we did it this way.

Q You mentioned you kept some numbers like help-desk-type calls or PCs that had to be wiped and rebuilt. Were you able to show some reductions in the malware problem after deploying FireEye?

A In terms of response, we were able to show some significant improvements in proactively detecting malware infections and addressing those things before any data theft occurred. While the volume of identified incidents had gone up due to increased visibility, our response time for infection-related incidents had gone down significantly. We were able to actively detect infections and take a more proactive approach to address malware infections.

The following statistics were observed in our environment following the implementation of FireEye and integration with the SIEM. (Please note, these statistics are just recorded numbers, while FireEye and SIEM implementation may contribute to the numbers, there are also other internal and external factors, such as incident response processes, targeted attacks, and time of the year that may also factor into the numbers):

- Comparing average data between the two months prior to the full implementation of FireEye and its integration to the SIEM, to the two months after the implementation in 2012, we have seen a 46% increase in detected network security incidents. A portion of this increase can be attributed to FireEye and the increased visibility associated with it.
- Comparing average data from before implementation 2012 a similar time frame in 2013, we actually saw a 35% reduction in network security incidents. I believe this is largely attributed to the actions we have taken to address detected issues, coupled with the finer tuning and scoping of our data sources.

Q Universities typically put a premium on sharing and openness. So they often worry about the privacy implications of inspecting traffic, and so on. Were you able to reconcile what you're doing with FireEye with those goals of the university?

A Correct. Privacy is definitely a huge concern, however at the same time, we have a long-established policy in the University that states: 'This is a Public Institution. All of your data and all of the activities conducted on our network and on our computers are covered under the Freedom of Information Act Law. If someone wants to see this data, they can.' We also try to drive the message that privacy protection and security work in tandem.

We do get requests from student media for justifications on why we block certain nodes. Our justification is: bad guys are using those channels to attack us. We don't block all of them, but do block the ones that are being used for active attacks. We are not sniffing people's traffic, but using these tools – FireEye and others – to help secure the University's data. I believe we have been effectively communicating that to the end-user community.

Q The desktop guys always are the ones dealing with malware. How does that work? FireEye tells you, 'We think this PC has this malware on it.' Whereas, the desktop AV is saying, 'Nope, everything's fine over here.' How did you cross that bridge to get something done?

A There were a couple of those cases where the desktop folks questioned us about the legitimacy of our reports. In those types of cases, one of the good things about FireEye is that it does allow us to see the actual p-caps [packet captures] of what's happening. Sometimes we'll actually even see the files that are changed along with individual registry entries that are added as FireEye is expanding these types of malware in its own VM [virtual machine]. We can then show the desktop support folks the data. Once they see the data, then they can go and verify that on the desktop or laptop itself to see whether there are any problems. One of the other good things about FireEye is that it has a very low rate of false positives. Although there were some doubts at the beginning, going through the 'forming, storming, and norming' stages, we were actually able to work out a pretty decent process where the desktop support folks were trusted us.

Q You mentioned you interfaced the FireEye units to your SIEM. What about where FireEye's telling you about malicious command and control addresses out there. Do you push those out to whatever you're using for URL blocking?

A Yes, we do. We block the IPs rather than the URLs. For a long time, we didn't have URL blocking capabilities, which is a shame to say, but we are now blocking the IPs. So, if it was things like Amazon or Google – obviously we were not going to block those - but for any type of malicious IPs that are reported, where there are more than three call backs associated with it, we would then triage those particular IPs and block them if necessary.

Once you put FireEye in place, you are going to have your eyes opened. You will end up seeing things that you would normally not see.

Q Are there any lessons learned you could pass on? Anything knowing what you know now you'd do differently?

A One of the big things is that once you put FireEye in place, you are going to have your eyes opened. You will end up seeing things that you would normally not see. Then you have to make sure you do have the capability—the processes and the capacity—to deal with the problems that are being detected. So I think we struggled with that a little bit up front. We didn't have a whole lot of processes worked out, especially with individual departments. At one point, we were notifying departments of potential infections, but the triage times were still not up to par. Eventually, we had to develop a process in place to say, 'if the reported infection is affecting one of our top-tier assets and the issue has not been resolved in 24 hours, then these are the actions that we must take.' I think aside from looking at the technology itself, having process in place and ensuring that you have the adequate capabilities and capacity is absolutely crucial to the success of the program.

Q You've been operational for a while now. How much resource and staffing does it take to run and use FireEye?

A I have one administrator who is primarily responsible for the management of the back-end interface, as well as the system itself. That person also manages our AV. It's not an entire FTE [full-time equivalent] managing the system. I have three separate incident responders. These incident responders do check FireEye, but primarily they go into our SIEM to get the correlated data and respond to incidents that way. They all have access to FireEye, but they don't regularly log onto it because FireEye does report back to the SIEM—that's a single pane of glass that we're using.

We were able to better detect different virus infections

Q Are there some things or features you've told FireEye you'd like to see added to the product?

A Getting some additional capabilities to really use FireEye as an IPS rather than IDS would be great.

Q On the GUI side, how about GUI and reporting? Are you happy with that?

A I think the GUI is actually very simple to use. It's probably one of the best GUIs that I have seen in products.

Q Now, do you do any custom rule creation, or you're pretty much just using what FireEye pushes out?

A We're pretty much using what FireEye is putting out. We do have a couple of custom rules in place, one of which is related to the ZeroAccess Trojan.

Q How have you found the support to be from FireEye?

A I think FireEye's support is excellent. So far we have not had any major support issues. Our folks have been able to get some really good response times from these folks.

SANS bottom line on FireEye Threat Prevention Platform:

- Anti-virus software and first generation intrusion detection systems are not effective in detecting advanced targeted threats.
- Universities have very active threat environments, and also have to deal with a high percentage of unamanaged or "sporadically" managed endpoints. Network-based advanced threat prevention is often the best starting point under these constraints.
- The FireEye Threat Prevention Platform showed a very low number of false negatives and false positives, reducing the staffing required.
- The use of the FireEye appliances increased the timeliness and number of security intrusions detected leading to reduced spread of compromise and decreased time and cost to mitigate.
- Integrating FireEye with a SIEM product enabled a 'single pane of glass' monitoring approach.
- To make effective use of products such as the FireEye Threat Prevention Platform, first: start by updating your incident detection and response processes, then look at network traffic volume and critical asset placement to optimize product selection, architecture and deployment.

FireEye protects the most valuable assets in the world from those who have them in their sights. Our combination of technology, intelligence, and expertise – reinforced with the most aggressive incident response team – helps eliminate the impact of security breaches. We find and stop attackers at every stage of an incursion. With FireEye, you'll detect attacks as they happen. You'll understand the risk these attacks pose to your most valued assets. And you'll have the resources to quickly respond and resolve security incidents. The FireEye Global Defense Community includes more than 2,700 customers across 67 countries, including over 157 of the Fortune 500. For more information, visit **www.fireeye.com**.