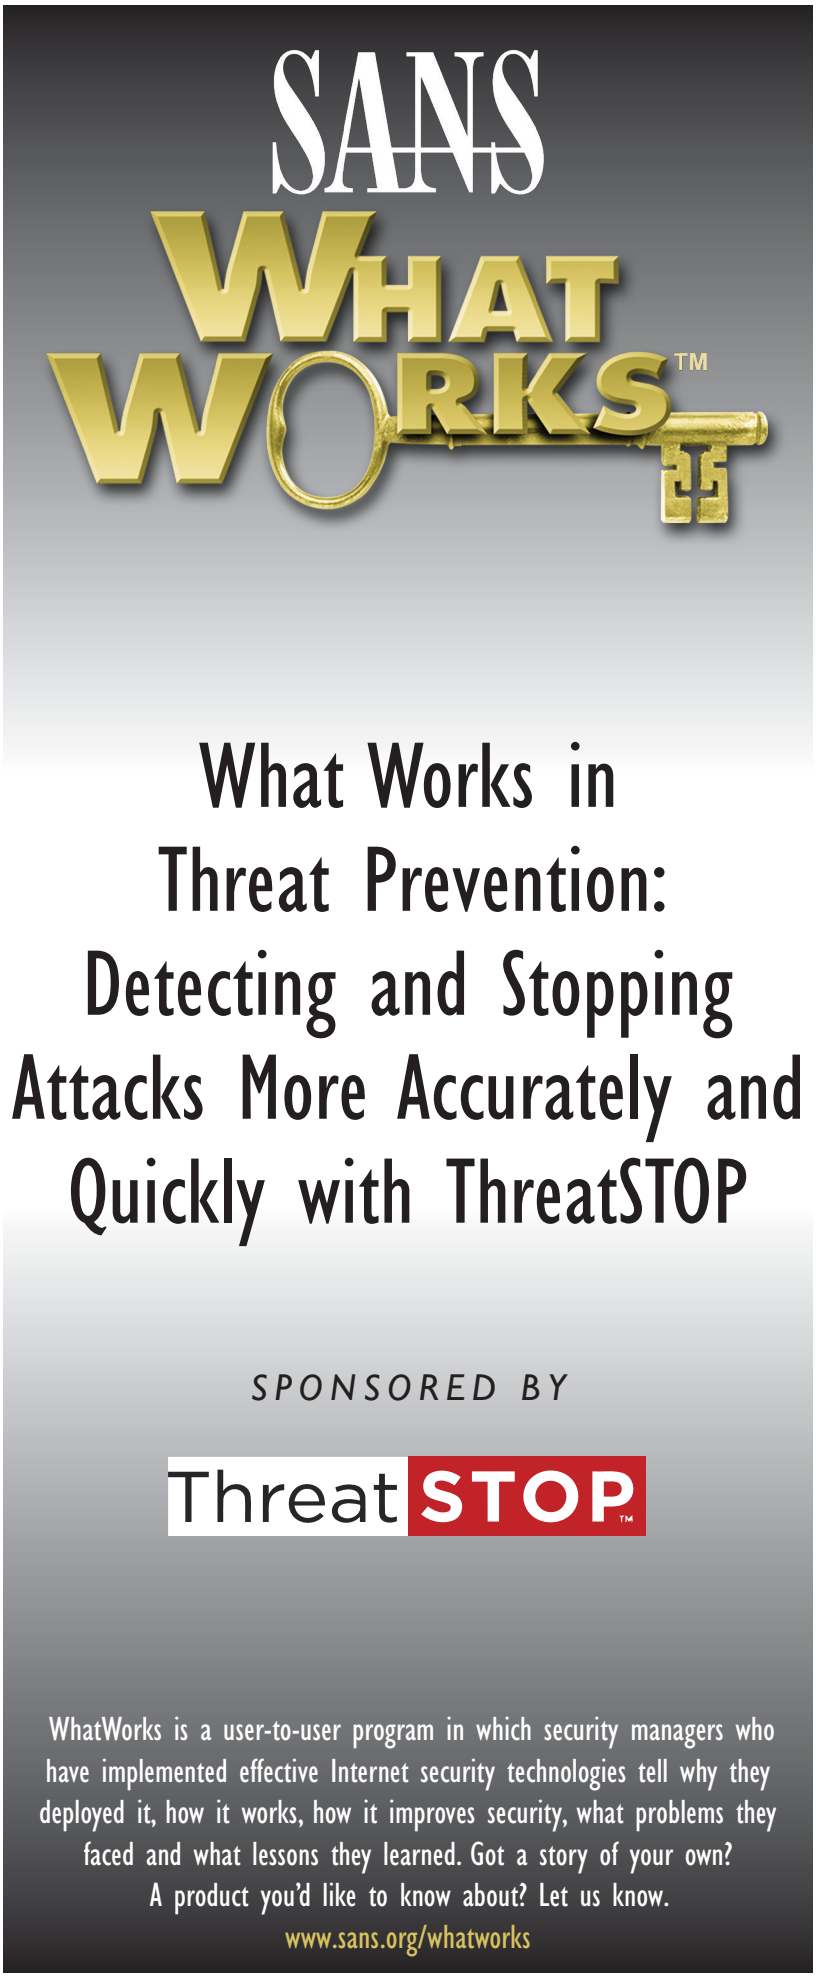
SANS
WHAT WORKS™
What Works in Threat Prevention: Detecting and Stopping Attacks More Accurately and Quickly with ThreatSTOP
SPONSORED BY
ThreatSTOP™
WhatWorks is a user-to-user program in which security managers who have implemented effective Internet security technologies tell why they deployed it, 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

SUMMARY

Press coverage tends to focus on breaches companies that have failed to protect their business systems and sensitive customer data. However, many enterprises have invested in improved processes, more advanced security products and threat-driven prioritization approaches to show immediate and measurable increases in both the effectiveness and the efficiency of their security programs.

ABOUT THE USER

Kenneth Compres is the Sr. Information Security and Integrations Engineer at Hillsborough Community College. Ken also serves as an associate professor of Cybersecurity at Mercy College and Bloomfield College in New York. At Bloomfield College, he serves as the subject matter expert for curriculum development for cybersecurity, information assurance and security programs. Mr. Compres developed the current disaster recovery plan for the Department of Veteran's Affairs in New York Harbor, New York. He has also developed several "code behind" for security monitoring at the VA Hospital in New York.

Mr. Compres has several research publications in IEEE (IT Governance, Compliance and Auditing Curriculum – A Pedagogical Perspective). He holds a Master of Science with a concentration in cybersecurity from Mercy College, and a Masters in Digital Forensics from the University of South Florida. He earned a Global Security Leadership Certificate (GSLC), and is a Fortinet Certified Network Security Administrator, Fortinet Certified Network Security Professional, Certified Ethical Hacker, Network+, Security+, A+, CCENT. He is currently pursuing a PhD in Computer Science from NOVA. He has been a guest speaker in New York's Security in the Workplace (2010, 2011), FireEye's Changing the Security Landscape, Orlando 2015, Information Security today. His most recent publication was published by Hillsborough Community College "Managing your digital footprint" October of 2012.

Q Tell me about Hillsborough Community College and your role there.

A I'm the Senior Security and Integration Engineer at Hillsborough Community College. We have approximately 40,000 students with approximately 4,000 employees, and close to 20,000 nodes; which does not include wireless guests, but actual devices that we control within the college.

Q Do you have a mix of residential dorm-type networks and academic networks?

A Yes, we do. We have a mixture of residential networks, BYOD devices, internal devices, and we also have guest networks that are specifically for people who are not students, employees, or contractors, but whom may be coming for a visit. So, we do provide access for those types of users.

...we also decided to bring in ThreatSTOP and put it at the forefront of our network so that they could help us detect these bots and these communications before they spread out.

Q Are you the equivalent of a CISO or do you report to a CISO?

A I'm the Acting CISO, along with the Director of Networks, whom I report to. We both fill that role.

Q What sort of problems were you having that led you to look around at solutions like ThreatSTOP?

A Interesting question. We first acquired ThreatSTOP, when it was in its infancy. It was still being pushed out. We looked at it, and we tested it. I remember one of the things that we discovered is that by using ThreatSTOP, we were able to mitigate a particular bot that was being spread out. One of the few places that didn't have that bug was the department that was using ThreatSTOP at the time. We looked at the tool and thought "why can't we implement this for the entire college – as opposed to only this particular department or building?" As it turned out, we switched it over, and put it forward for the entire college. It completely helped our dynamics on threat defense.

We looked at the tool and thought "why can't we implement this for the entire college – as opposed to only this particular department or building?"

Q So, a particular department had been trying out ThreatSTOP, and then you noticed that they were the only area where bots had not been detected, correct?

A Correct. That was the only area where botnets and malware had not spread to.

Q How had they deployed ThreatSTOP in that department?

A At that point in time, every department and every campus in the college had its own set of threat mitigation tools; keep in mind that we're a very large community college. Most of the time, when you think about community college, you think of single campus, single city and single exit and entrance point. We have seven campuses throughout the Tampa Bay area, including South Shore and Lakeland which are 26 miles away from the main campus, but they're all campuses of Hillsborough Community College. At that time, we had campuses that were just using ThreatSTOP, but they also had to gain access by coming to our central location. Nevertheless, they also had their own equipment. So, what we did, as the years went by, was centralize all of the communication between the campuses and – within the campuses and our central location – so that we could have a more granular control of the communications. With that being said, we also decided to bring in ThreatSTOP and put it at the forefront of our network so that they could help us detect these bots and these communications before they spread out.

Q ThreatSTOP essentially feeds IP addresses into firewall access policies, and it also has some other more direct DNS-type firewall function. Which of their products or services are you using?

A We have ThreatSTOP Shield and ThreatSTOP DNS Firewall, At that time we had a CISCO ASA, and the integration was seamless. When ThreatSTOP detects a particular communication going out into a known botnet IP, it would automatically detect/inject that block into our firewall so that we lose no time in detecting it. One of the things I have configured in our ThreatSTOP is the ability to get our monthly reporting so that we can see the dynamics and how many systems have these types of communications over the course of the month, and monitor that on a continual basis.

Q If something gets on the inside, or if some user on the inside goes to resolve a DNS address and ThreatSTOP knows that it's malicious or shady, you're essentially stopping it at the DNS resolution side. But, if there's some communication that's going through the firewall to certain IP addresses and you may not have made a DNS call, then you would be able to block that at the firewall with that deny rule?

A Absolutely correct.

Q It was originally used in one place. You saw benefit from it. But, then you had to expand. How is it deployed physically? When you wanted to expand to get wider coverage, what did you have to do?

A We originally had a Cisco firewall that took care of all the communications inbound and outbound, and ThreatSTOP was placed behind the firewall. We had the firewall send all communications to the ThreatSTOP system, which would digest the information and create its ruleset based on the information given. Prior to the college-wide deployment of ThreatSTOP, we had every campus with its own ingress/egress rules. We simply setup the rules to be seen by ThreatSTOP, campus by campus.

Q As part of your consolidating all the internet communications, were you able to make sure everything was flowing through that firewall that had the ThreatSTOP feed?

A Correct.

Q How long have you been using ThreatSTOP?

A We have had it now for about seven years.

Q Since you've been using it for a while, was it called ThreatSTOP when you got started?

A I don't believe it was called ThreatSTOP when we got started – it was just an idea.

Q Over those years, did you look at any other solutions? Did you have to do any tradeoffs or evaluations of other things?

A We evaluated a couple other solutions, however, when you consider the cost of the competition (FireEye) vs the benefits, it simply made sense to go with ThreatSTOP, the reality is that we felt ThreatSTOP met and exceeded our expectations, and we really liked the fact that it was easy to use, easy to maintain, and easy it implement. The management interface was very, very intuitive, with a very low learning curve. It also helped that they listened to our suggestions. Whenever we needed a particular change in the view or in the system itself, or were doing an upgrade; it was actually very easy to reach out to the engineers at ThreatSTOP who responded and actually worked with us and made it happen.

Q Physically speaking, what you're running from ThreatSTOP is software you put on a virtual server?

A When we originally built it, it was a physical server. As time went on, we moved the system to a virtual and it has been a virtual ever since.

Q When you look at those types of solutions, there is the false positive problem. If it's not knowledgeable about certain IPs or URLs that might be bad, it lets things through. The other is the false negative side where it might have blocked the legitimate DNS request. Have you seen false negatives and false positives?

A I have seen very few false positives. Whenever we get a complaint about a particular IP being blocked, and we've analyzed the IP, it turns out to be a real threat. I would say it has been less than 1 percent over the years. Considering we have approximately over 350,000 session requests at any given moment, it's been very minimal.

ThreatSTOP is a solution that any company that's serious about their security needs to invest in.

Q Is there tuning you need to do, either when you first got started years ago or as you've been going?

A Originally, we had allowed communication – incoming and outgoing – anywhere, and we only used ThreatSTOP to mitigate attacks. As time went on, I was able to move forward, blocking countries where do don't do any business, and ThreatSTOP was essential in making that transition easy. ThreatSTOP is a solution that any company that's serious about their security needs to invest in.

Q How do you administer it? What is the staffing load?

A It's very minimal. If you consider that the person who handles the firewall has very little interaction with the device itself, except to apply updates every once in a while, or if there's an update to the firewall itself, you want to make sure that the version is compatible with the current appliance. We haven't experienced many problems with that. It's very automated. It's very self-managed. The person that runs the firewall will actually find their jobs easier with this service added onto their network, and your technicians will appreciate the reduced load of malware in the environment.

Q Do you run it essentially automated?

A Yes. Once it detects a system attempting to communicate with a known command control center, it will implement – inject – that block rule onto our firewall. We do get weekly reports that I have my team review, and whenever there is a specific attack, we take action. These alerts are sent to our technicians with information as to where they originated from and the type of system making the calls.

Q Sometimes, the legitimate site will be compromised and serve as a command-and-control center for a while, and now you have a role blocking access to part of a legitimate site that, when it gets fixed, you don't need to block any more. Do they also push out the removal type to remove the deny?

A Yes, they do. There has been one occurrence over the period of time where I had to manually remove a particular IP from the list, but that rarely happens.

Q From a latency and speed and performance point of view, have you seen any problems?

A Absolutely not. I think that, if you have the right equipment in place, you will actually see an improvement on your communications because you don't have as many devices bogging down your network with malware and other type of coordinated attacks.

It's very automated. It's very self-managed. The person that runs the firewall will actually find their jobs easier with this service added onto their network, and your technicians will appreciate the reduced load of malware in the environment.

Q How do you know if ThreatSTOP starts pushing out deny rules, and you're effectively blocking something? Are there blocking alerts from the firewall / log events from ThreatSTOP? How do you then, from that information you're getting from ThreatSTOP or the firewall, figure out which machines you have to go look at that are making these calls?

A ThreatSTOP tells you on the reports or in the logs online. You can actually see the source and destination of the device making the call to the C&C center. You can set up your firewall to alert you when a particular new device that is considered high priority has been detected and has been blocked. I would normally have one of the technicians go out and scan it and take a closer look at it and see what's going on with that particular system.

Q You're essentially going by the source IP. Is there any other information they give to you; if it's a BHCP address or if you have NAT on your internal networks anywhere?

A When ThreatSTOP sees a call to a C&C center, it sees the internal IP address – since we have an internal appliance for it. The firewall gives it the source IP, the NAT IP, and that allows us to narrow down whether this is a wireless user or it's an internal computer. For the most part, if it is a wireless guest user or a wireless student user, we really don't look at those because since they are not intrinsically part of our network; they are isolated from our network. Nevertheless, we try to provide the portion of the protection to them, too. However, when it's an internal user, the firewall provides the source IP address of this device, the destination IP address and whether this was a DNS request and what type of communication was performed.

Q You mentioned you had some suggestions/ requirements which ThreatSTOP was responsive to. What sort of features have you asked them to add?

A I have talked to them about ThreatStop doing PCAP with an emphasis on the communication detected.

Q Is there anything on your wish list or things you hope to see the product do in the future?

A Yes. I would like to see ThreatSTOP enable us to integrate the system with ForeScout or a Network Access control device so that it can automatically alert the device and have the NAC take action and block the node in the network. I have been working on code to make this work; however, if ThreatSTOP could facilitate the solution that would be great.

Q How have you found the support from ThreatSTOP?

A The support from ThreatSTOP is excellent. I have had no problems with their support which is one of the things I've really enjoyed. When we want a particular change within the system, a particular look for a report or updating one of our systems, we are able to reach out to them and they are very responsive with it.

Q I assume, every year, you have to renew the procurement. Are there any metrics you collect or numbers you have that you present to management to validate why you'd want to keep this solution?

A Yes. We collect details of inbound and outbound attacks. For example, we completely block communication with China because we don't do business with China. We are able to show them our collective deny-service attacks being launched from this particular country or active nation, how much has been blocked, how well it's functioning, and what it would have cost us had these attacks been successful. They were able to see the value in it. Once the companies see the numbers and how those blocks are occurring before the attack is fully affected, they will start to understand the real value of it.

Once the companies see the numbers and how those blocks are occurring before the attack is fully affected, they will start to understand the real value of it.

Q How are you doing DNS? What are you using for DNS resolution?

A We have internal and external DNS servers. These servers are managed by us, but I can see in the future, moving them to a cloud-based system.

Q From a performance point of view, have there been any issues?

A No issues. As a matter of fact, I believe it's helped improve our system's performance. With regards to the reporting, I think people will really enjoy being able to see the data, the attacks being mitigated and the fact that their systems have fewer outages, and that in itself speaks volumes. When you're able to demonstrate how this particular system is protecting your network, reducing the amount of denial-of-service attacks, reducing the cost of troubleshooting specific issues and actually preventing issues before they become a real problem, I think companies will appreciate the value of ThreatSTOP.

Bottom Line:

A community college found that ThreatSTOP enabled them to detect in-process targeted attacks more quickly and with few false positives. This enabled them to reduce the number of resources reached by attack payloads and significantly reduce the business impact of those attacks. Integration with existing security controls and the accuracy of the threat intelligence supported an increase in the level of security without requiring increased staffing or investment in additional security products.