SANS

Using WhiteHat Sentinel Dynamic and Static Solutions to Increase Application Security Before and After Production Deployment

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

SUMMARY

While attackers continue to exploit vulnerable websites, many enterprises have been able to greatly increase the security of those applications and the maturity of their overall software development processes and avoid such breaches. This Case Study details an enterprise's success in using WhiteHat Sentinel to more rapidly and accurately find and fix vulnerabilities in web applications. The user also provides details on how his company evolved to avoid operational vulnerabilities by using WhiteHat Sentinel to test pre-production code before approving it for release.

ABOUT THE USER

The user interviewed for this case study has requested anonymity to maintain confidentiality, but has allowed us to refer to him as the CISO and VP of an emerging technology company. The SANS WhatWorks program can help our security community at large make more informed decisions by encouraging seasoned professionals from major user organizations to share their stories without revealing the name of the organization.

ABOUT THE INTERVIEWER

John Pescatore, SANS Director of Emerging Security Trends
Mr. Pescatore joined SANS in January 2013 with 35 years' experience in computer, network and information security. He was Gartner's lead security analyst for 13 years, working with global 5000 corporations and major technology and service providers. Prior to joining Gartner Inc. in 1999, Mr. Pescatore was Senior Consultant for Entrust Technologies and Trusted Information Systems, where he started, grew and managed security consulting groups focusing on firewalls, network security, encryption and Public Key Infrastructures. Prior to that, Mr. Pescatore spent 11 years with GTE developing secure computing and telecommunications systems.

Mr. Pescatore began his career at the National Security Agency, where he designed secure voice systems, and the United States Secret Service, where he developed secure communications and surveillance systems. He holds a Bachelor's degree in Electrical Engineering from the University of Connecticut and is a NSA Certified Cryptologic Engineer. He is an Extra class amateur radio operator, callsign K3TN

Q Tell us a little bit about yourself and the role you play at your organization.

A I'm the CISO and VP of IT for an emerging technology company. My job has two parts – to ensure the infrastructure is locked down and secure as well as the product.

Q Is your company's product a software product?

A Yes.

Q What problems persuaded you to start looking at solutions like WhiteHat?

A It was the traditional issues: Do more with less, figure out where vulnerabilities may reside, and ensure that your information security program has application security built into it.

Q Were you using any application security products and/or processes that you had in place already?

A Yes. There were tools in place in a number of areas, and what I felt was missing was the 24/7 continuous scanning and monitoring of the application. So that's when we pursued WhiteHat.

Q There was an online web exposed application, and you wanted to have more continual monitoring of vulnerabilities in the app?

A Correct. We started with the Internet facing application, and then moved upstream towards the pre-production phases and scanning for vulnerabilities then. Using WhiteHat was straightforward, and enabled us to look at what was in production first, and then later on pre-production. Vulnerability scanning during pre-production actually took a shorter period of time to get running. It was powerful to have that time-to-value in such a short window – literally, within hours.

Q When you started out, how did you go about looking at solutions? Did you look at a variety of solutions – did you do bake-offs or RFPs?

A There were a number of vendors that were contacted, but most of them couldn't perform in the accelerated timeframe that we needed. We needed vendors to start immediately, and most could not. We ended up choosing WhiteHat because they stood up the solution in a matter of hours. The other vendors – the services companies, the professional services organizations - only provided a snapshot in time, providing a pen test or code analysis once or twice a year, but we needed something that had more scanning capabilities in real time. The other scanning companies we contacted had some issues with the comprehensive view of the application. They were limiting themselves to the scans – but what they weren't performing were business logic tests. Business logic says the application doing A, B, and C and really shouldn't do anything else. A cyber criminal comes in and realizes that the application can do A, B, C, D, E, and F. WhiteHat was able to demonstrate how they were going to perform those business logic tests, and ensure that the entire application was looked at from beginning to end in that capacity. Other scanning vendors weren't able to scan applications 24/7 that we needed. We also saw way too many false positives in other offerings. WhiteHat has the validation if the vulnerability that's discovered is a false positive or not, and that's extremely powerful for us since we're a small group focusing on building the software securely. When you have WhiteHat as a strategic partner in the solution, it's easier to figure out. That becomes critical in our environment because we don't have the luxury of an endless budget to fund all the research required with other scan vendors. For example, vendor X would say "we have a high percentage hit rate for accuracy." However, when we look at the results, over 20 percent are false positives. That becomes problematic and time consuming.

We ended up choosing WhiteHat because they stood up the solution in a matter of hours.

When you have WhiteHat as a strategic partner in the solution, it's easier to figure out.

Q Give some detail on what WhiteHat does to validate results to eliminate false positives before they're passed on to you?

A WhiteHat has an engine called Sentinel, which allows an organization to scan their environment. When an environment is scanned with WhiteHat, that data is looked at and reviewed by subject-matter experts, and those application security subject matter experts then validate for you that the scans are truly vulnerabilities. If the results are false positives, they get put to the side, and when you're presented with your data, you have a dashboard with meaningful data and metrics with the result of that scan. The false positives are weeded out, and you have the real vulnerabilities presented to you.

Q There's human expertise in the loop prior to the results going to you?

A Correct – this is critical for a number of reasons – primarily, having a subject-matter expert as an extension of your team validating that application.

Q Give an idea of the scale of this application and the scanning that's done. Is there both straightforward scanning and credentialed pen tests?

A The application has an Internet facing component, and a back end component; multiple components make up one application. There's the Internet facing component of it, but then there's also the authenticated scans that have to be performed. What's nice about the way WhiteHat does it, is they can scan the application as an Internet facing component; but then you also have the authenticated scans where they could look inside to say what happens when a user authenticates, and it simulates what a user is doing. That's powerful for a number of reasons. For one, what happens after you authenticate to that application; are there escalated privileges that could be compromised? Also, does the application perform as designed after the person has authenticated into the environment?

Q What sort of vulnerabilities has this type of testing brought out or exposed?

A We're an agile development shop, so we're releasing code quite often, so it's beneficial to have 24/7 scanning. The types of vulnerabilities that we've seen have been the traditional cross-site scripting/cross-site request forgeries. The types that, in my opinion, have been the low hanging fruit that are validated using a solution like WhiteHat. They are easy to go back and resolve. We had some internal training to show the developers why you create a certain process or framework to eliminate those cross-site scripting vulnerabilities.

Q How does the workflow happen, being that you're an agile shop? Does the scanning go asynchronously and keep happening? And if there are results from WhiteHat, do they go to you? How does that information flow from the security side to the development side to get things fixed?

A There's a framework. We try to make sure security is built in – there have been discussions during the scrum meeting about authentication, accessibility in emphasizing security as part of the development cycle there. That iterates every time there's new features or functionality built in so it's caught up front. Then as they're going to their sprint, there's another crosscheck with security. Then as the code is released, we have a platform where WhiteHat can scan that platform continuously. So, it's been pretty straightforward. In my opinion, it's very forward thinking because we're able to catch any of the security discussions or topics early during the architecture and tech design phase; and then also because of the continuous scanning, we're able to capture any potential vulnerabilities in pre-production and production.

Because of the continuous scanning, we're able to capture any potential vulnerabilities in pre-production and production.

Q When WhiteHat results come, are they integrated into some trouble ticket system? How does that information flow?

A As of now, we're manually looking at those, but we're migrating to a new ticketing system, and we're going to be using the WhiteHat API as part of the ticketing system so it automates the feeds into work flow system.

Q You mentioned scanning going on in the Internet facing production app. Is there also testing done on final QA or in the final pre-production stage?

A It is my experience in the industry that when somebody tells you that pre-production and production have the same code base and the same systems and infrastructure, chances are they're not the same. We have WhiteHat reviewing pre-production and production environments to ensure that those applications are truly configured the same way, making sure that any of the vulnerabilities are caught in both environments before they're released.

Q Are your pre-production and production environments hosted locally on your own servers? Are they out at some cloud service provider?

A They're all internal at this time.

Q Is there some way you can give WhiteHat access to the pre-production side in a secure way?

A Yes, they have a VM that runs, and what we do is use the WhiteHat virtualized install that feeds information to WhiteHat about the vulnerabilities. One of the nice things about that is it's not taking source code, it's not taking binaries, it's actually taking the results of the scan and feeding them up to their portal, which is critical for us because we don't want any of our intellectual property or any of our source code or binaries leaving our environment. We want to keep everything on premise and not have a risk of releasing any source code out into the cloud or offsite.

Q WhiteHat has features where you can take the results of a scan and feed into web app firewalls to provide protection before the fixes are made. Are you doing anything like that?

A We will be evaluating that feature in Q3 to figure out how and when we want to integrate.

Q When you made the decision to go with WhiteHat, how long from that point was it before you were up and running?

A The Internet site was less than one hour that just consisted of providing the URL and the credentials for an admin user, user one and user two, so that they could test those environments. The pre-production environment took less than one hour, and that was because we had to procure the server, set up the server, install and then test. We had a great experience with both of those environments. We were assigned an expert from the Threat Research Center (TRC) from WhiteHat who walked us through the installation and got everything up and running.

Q What sort of staffing level does it require you to deal with the whole WhiteHat side of the process?

A For staffing, there's one security engineer who spends approximately 50 percent of their time reviewing and then walking through the application security vulnerabilities discovered by WhiteHat. This has actually been reduced since we've implemented the WhiteHat solution. We've seen the number of vulnerabilities reduced to a very low number now, which has been very powerful and very positive. I would say it's a half-time position for one person.

We've seen the number of vulnerabilities reduced to a very low number now, which has been very powerful and very positive.

Q What sort of metrics have you been capturing to convince management, that the money they're spending on application security, WhiteHat in particular, are leading to an increase in security? You mentioned a reduction in vulnerabilities per line of code or per time.

A We had to go back and review what it would take to hire a full time employee, as well as purchase additional software we would need to support the application scanning efforts. And the issue is, to find a qualified, trained person in application security is almost impossible right now – and there are no signs of it slowing down. So, it's either staff augmentation or find a solution like WhiteHat. And for us, it made business sense to go down the path of WhiteHat because you're lining up these subject matter experts in the Threat Research Center that are trained, qualified personnel that are going to perform these types of tests non-intrusively on your environment. So, when I went to our leadership team to explain why we needed this, what we were going to do, and how we were going to do it, it just made business sense that to go down this path than hire the FTE, purchase the equipment, purchase the software and have that one FTE running this program 24/7/365.

Q On the TCO side, are you able to show management what you're spending on WhiteHat is less per year than what it would take do it yourselves? Are you able to show this reduction in vulnerabilities over time, as well?

A Correct. The crazy part is, when you look at what it takes to on-board a full-time person versus the application and the number of apps that you need to test, the WhiteHat solution is so much more affordable. It's very powerful when you have to build your TCO around that.

Q What has your experience been with support service from WhiteHat and access to those subject-matter experts?

A Our experience has been very positive. Within the WhiteHat product portal, you have the ability to ask questions about the identified vulnerability. If more clarity is needed around that vulnerability, you can set up a call and walk through it with one of the engineers from the TRC. The portal also has a place for you to ask questions and get answers, so you have customer support there.

Q How long have you been operational now with WhiteHat?

A Over 14 months.

Q Over that time, are there any lessons you learned? Any things that, knowing what you know now, you would have done differently?

A I would say it would start earlier in the process, in order to plan and automate.

Q Are there any wish-list items or things you've passed on to WhiteHat you'd like to see added to the products and the services?

A Yes. WhiteHat has a scoring module called the WhiteHat Security Index (WSI) – something we're really excited about because it's a security index that shows you how you are doing as a company and your vertical – then you and your organization across the universe of WhiteHat clients. We're looking forward to receiving some reports that we've asked for as well.

Q Is there anything else I didn't ask that you'd like to mention?

A The key thing for us, if I had to summarize, is time-to-value. Less than four hours, up and running, less than four hours for pre-production, less than an hour for production. And showing that value in less than an hour is powerful when you're working with an organization that's moving fast. Then on the support side, I think it's critical to stress that the TRC is comprised of application security experts. I would also like to mention that when someone gets hung up, or an engineer has a question, they can actually interact with that individual at the TRC very quickly to get the results that they need or the questions answered they need. Very powerful all around.

Q Does your company do anything to judge WhiteHat's performance? Are they missing vulnerabilities? Are you getting reports from external security researchers that are validating that they're not missing things?

A We have a number of application security alerts that we collect from vendors and system providers. What's interesting about this is, WhiteHat's often already posting this information in a number of areas. By the time we receive application security information from other email lists or product updates, WhiteHat already has a posting in their environments – the website, blog, or email, so it's nice to have that cross check. We also perform pen testing twice a year in our environment, so it's nice to have the validation that WhiteHat isn't missing anything. Once again, very powerful.