IT Central Station
Unbiased reviews from the tech community

Case Study

Vectra AI

Mark Davies

Security Operations Specialist at a tech services company with 1,001-5,000 employees

Review by a Real User
Verified by IT Central Station

What is our primary use case?

We use Vectra AI to sniff the network using Ixia taps so that we can identify potentially malicious activity on the network and at all points of the kill chain. What it's exceptionally good at doing is correlating seemingly unrelated events.

It's in our data center, but the versioning is controlled by Vectra. They push it out discreetly so I don't have any touch on that.

How has it helped my organization?

We have 89,000 concurrent IPS that we're analyzing and it's distilled it down to under 1,000 IP addresses that warrant deeper investigation. It's filtering out 99 percent of the traffic that would otherwise be noise, noise that we would never get through.

The solution captures network metadata at scale and enriches it with security information, but that's because we are using the API calls to inject our CMDB data into the brain. It speeds things up quite significantly. Being an enterprise, sometimes it can take a day or two just to find the person responsible for looking after a particular server or service. This way, the information is right there at our fingertips. When we open up the GUI, if we have a detection we look at the detection and see the server belongs to so-and-so. We can reach out to that party directly if we need to. It streamlines the investigation process by having the data readily available to us and current. Each one is unique, but typically, from initial detection to completion of validation (that it's innocuous or that there's something else is going on) it's within 24 to 48 hours

It also provides visibility into behaviors across the full lifecycle of an attack in our network, beyond just internet gateway. It gives us visibility for when something is inside the network and it's maybe doing a lateral movement that it wouldn't normally be doing. Or if we have a system that has suddenly popped up on the network and we can see that it's a wireless router, for example, we pick that up right away. We can see it and we can deal with it. If people put unauthorized devices on the network — a wireless router from home — we can pick that up right away and deal with it.

In addition, Vectra triages threats and correlates them with compromised host devices. We can do a search based on the threat type and get the host. It streamlines things and makes it faster to get to the root cause of an issue.

And while it hasn't reduced the security analyst workload in our company, it has reduced the workload in that analysts are not having to look at stuff that absolutely means nothing. There is still a lot to do, but it has allowed us to focus better on the workload that needs to be done.

It has also increased our security efficiency. It has reduced the time it takes us to respond to attacks by 100 percent. If you're not aware of it you can't respond to it. Now, it's making us aware of it so we can respond to it, which is a 100 percent improvement.

The solution enables us to answer investigative questions that other solutions are unable to address. We will detect the fact that there is some suspicious domain activity going on — a DNS query is going out to MGAs and it really shouldn't be. The other systems are just passing that through, not even realizing that it shouldn't be happening. We see them and we can take action on them.

What is most valuable?

The dashboard gives us a scoring system that allows prioritization of detections that need attention. We may not necessarily be so concerned about any single detection type, or event, but when we see any botnet detections or a brute force attack detections, we really want to get on top of those.

What needs improvement?

The solution's ability to reduce false positives wasn't very good, initially, because it was picking up so much information. It took the investment of some time and effort on our part to get the triage filters in place in such a fashion that it was filtering out the noise. Once we got to that point, then there was definitely value in time-savings and in percolating up the high-risk events that we need to be paying attention to.

I'd like to be able to get granular reports and to be able to output them into formats that are customizable and more useful. The reporting GUI is lacking.

For how long have I used the solution?

I've been using Vectra for three years.

What do I think about the stability of the solution?

The stability is excellent.

What do I think about the scalability of the solution?

We've had no issues so far with the scalability. Right now, it covers about 90 percent of our network. We are considering increasing the usage to incorporate it in the new cloud environments that we're standing up.

How are customer service and technical support?

Their technical support is excellent.

How was the initial setup?

I was not involved in the initial setup, but I was involved in a review of the setup when I took it over, to make sure that it is doing what it's supposed to be doing. The initial setup would have been straightforward, but it would have been very large.

The implementation strategy would have been to make sure that it got to all the places that it needed to be, and to work out a way to make that happen by getting the Ixia taps into the right locations in our enterprise.

In terms of staff from our side involved in deployment, it's web-based so there weren't a lot. Maintenance is ongoing from Vectra and they do it on the back-end. It just works. It's a black box for us.

What other advice do I have?

Take time to understand how the triage filtering works and standardize it early on. Use a standardized naming convention and be consistent.

It's a very effective tool, but if you don't pay attention to what it's telling you, then it's like anything else. If you don't use it, then it's no good. You have to trust that what it's telling you is correct and then you can take the appropriate action.

For the most part, the users who log into it in our company are people on the security operations team. It's pretty much a closed tool. Access is limited to the people in the security center of excellence.

In terms of the solution's ability to reduce alerts by rolling up numerous alerts to create a single incident or campaign for investigation, we don't use it that way. We've set up enough triage filters over the course of the last year-and-a-half to get all the noise out of the way; stuff that is either innocuous or really isn't bad. Then we're

focusing on what's left, which is typically, for lack of a better term, the bad stuff or the stuff that we need to pay attention to.

Regarding the solution's privileged account analytics for detecting issues with privileged accounts, we've used it, but not to the extent that we would like to. We just don't have enough manpower to be able to do that at this point. But it's important because we can see when an account is doing something that it shouldn't be doing, or that it doesn't normally do, or that it's connecting to a place that it doesn't normally connect to, or that it's escalating its privileges unexpectedly. We see all that and then we can respond accordingly.

Which deployment model are you using for this solution?

On-premises