IT Central Station
Unbiased reviews from the tech community

Case Study

Vectra AI

reviewer1263180

Cyber Security Analyst at a financial services firm with 1,001-5,000 employees

Review by a Real User
Verified by IT Central Station

What is our primary use case?

The original use case was because we had some legacy stuff that doesn't do encryption at rest. Compliancy-wise, we had to put in some additional mitigating actions to protect it. That was the start of it. Then, we extended it to check other devices/servers within our network as well.

We are on the latest version.

How has it helped my organization?

It is doing some artificial intelligence. If it sees a server doing a lot of things, then it will assume that is normal. So, it is looking for anomalous behavior, things that are out of context which helps us reduce time. Therefore, we don't have to look in all the logs. We just wait for Vectra to say, "This one is behaving strange," then we can investigate that part.

We have implemented it fully now. We have done some training and filtering on it. Now, every alert that we see means that we need to investigate. It sees roughly 300 events a day. The majority are normal behavior for our company. So, there are about 10 to 15 events a day that we need to investigate.

The solution triages threats and correlates them with compromised host devices. It looks at a certain IP address, and if you're doing something strange, then it will give us an alert. E.g., normally John Doe is logged into it for four days, going to server XYZ. If all of a sudden, it's in a different timescale, going to server B, then it will send us an alert.

We have privileged accounts. They have a

specific names, and if I see those names, then I investigate a bit more thoroughly. That's our policy. I don't know whether Vectra does anything different with them.

The solution gives us more tickets. If we did not have Vectra, we wouldn't have those tickets. So, it's actually increasing them. However, it is improving our security with a minimum amount of work. That's the whole purpose of the device. We have 10 to 15 events that we need to look into a day, and that is doable.

The solution creates more work for us, but it is work that we are supposed to do. We need more FTEs because we need more security.

What is most valuable?

We mainly use it for the detection types, checking dark IPS or command-and-control traffic.

We bought Recall so we can have more information. Recall is an addition onto Vectra. We haven't enabled Recall yet, but we will. So, if there is an incident, we can investigate it a bit further with Vectra devices before going into other tools and servers. This gives us the metadata for network traffic. So, if we have a detection, we can check with Recall what other traffic we are seeing from that device, if there is anything else. It's mainly a quick and dirty way of looking at it and getting some extra information to see whether it's malicious.

We found that the solution captures network metadata at scale and enriches it with security information. This is one of the reasons why we added Recall, so the alert gives us information on where we need to look, then we can investigate a bit further. For example, a certain device is sending data to command-and-control server, then we can investigate whether that is really happening or just a false alarm with the metadata in Recall. It makes it easier to find out.

What needs improvement?

We would like to see more information with the syslogs. The syslogs that they send to our SIEM are a bit short compared to what you can see. It would be helpful if they send us more data that we can incorporate into our SIEM, then can correlate with other events. We have mentioned this to Vectra.

It does some things that I find strange, which might be the artificial intelligence. E.g., sometimes you have a username for a device, then it makes another. It detects the same device with another name, and that's strange behavior. This is one of the things that we have with Vectra support at the moment, because the solution is seeing the device twice.

For how long have I used the solution?

We started the pilot roughly a year ago. So, we started small with a pilot on part of the systems, then with two other vendors. Afterwards, we decided to buy it.

Now, it's almost in production. It's still a project in the end phase, as we are still implementing it. But, most of it has been running for a year.

What do I think about the stability of the solution?

So far, the stability has been good. There are no issues. It's never been down. It has been updating automatically on a regular basis and there are no issues with that where it has stopped working.

One person will be responsible for the deployment, maintenance, and physical upkeep; a person from the service delivery team will keep the device up and running. The security analysts (my team) deal with the alerts and filtering.

What do I think about the scalability of the solution?

The part that we designed is not really scalable. They have options, and there is some room for improvement. If we need to scale up, which we have no intention of doing, then the physical devices need to be swapped over for a bigger one. Other than that, we have some leeway. This came up in the design with, "What are your requirements?" and those requirements have been met, so that's fine. They will probably be met for the foreseeable future.

At the moment, we don't have Tier 1 and Tier 2. Instead, we have a small team who does everything. I am mostly using it. There will be three security analysts. Then, we have a number of information security officers (ISOs) who will have a read-only role, where they can see alerts to keep an eye on them, if they want, and be able to view the logging and see if they need more information. But, there are three people who will be working with Vectra alerts.

How are customer service and technical support?

We are in contact with the Vectra service desk. If you send them ideas, they talk about them and see if they can incorporate them.

Which solution did I use previously and why did I switch?

We decided that we wanted to have an alert within 30 minutes, which is doable with this solution. It fulfills our needs. However, we didn't have this before, so it has increased our time, but for things we need to do.

How was the initial setup?

The initial setup is relatively straightforward. They have security on a high level. There are a lot of logins with passwords and very long passwords. This made it go a bit longer. However, the implementation is relatively easy compared to other devices.

We made a design. That's what we implemented.

What about the implementation team?

Initially, it was set up in conjunction with Vectra. When we put it into production, the majority was done by me, then checked by a Vectra engineer. If I had issues, I just contacted Vectra support and they guided me through the rest of it.

The Vectra team is nice and helpful. The service desk is fast. They know what they are doing, so I have no complaints on that part. We have a customer service person who knows about our environment and can ask in-depth questions. He came over as well for the implementation to check it, and that was fine. The work was well done.

What was our ROI?

The solution has reduced the time it takes us to respond to attacks. It sends an email to our SIEM solution. From that SIEM solution, we get emails and tickets. Therefore, the time between an alert coming up and a ticket is reduced. This is for tickets that we monitor regularly. Within 15 to 20 minutes, it gives us an alert for the things that we want. Thus, it has greatly reduced our measurable baseline.

The return of investment is we have tested it so sometimes we have auditors who do pen tests and see them. That's the goal. It seems to be working. We haven't found any actual hackers yet, so I'm not completely a 100 percent certain. However, we found auditors who are trying to do pen tests, which essentially the same thing.

What's my experience with pricing, setup cost, and licensing?

The license is based on the concurrent IP addresses that it's investigating. We have 9,800 to 10,000 IP addresses.

There are additional features that can be purchased in addition to the standard licensing fee, such as Cognito Recall and Stream. We have purchased these, but have not implemented them yet. They are part of the licensing agreement.

Which other solutions did I evaluate?

We investigated Darktrace, Vectra, and Cisco Stealthwatch.

Darktrace and Vectra plus Recall were similar in my opinion. Darktrace was a bit more expensive and complex. Vectra has a very nice, clean web GUI. It easier to understand and cheaper, which is one of the main reasons why we chose Vectra over Darktrace.

Darktrace and Vectra are very different, but eventually for what we wanted it to do, they

almost did the same thing. Because Darktrace was a bit more expensive, it was a financial decision in the end.

I did the comparison between Darktrace and Vectra. They did almost the same thing. Sometimes, there are differences that Darktrace did detect and Vectra didn't. For the majority, we didn't find any actual hackers. So, it's all false positives, eventually. Both of them are very similar. The big thing is the hacker activity. They both detected it in the same way. But, in the details, they were different.

The options for Stealthwatch were a bit limited in our opinion for what we wanted it to do. Stealthwatch is network data, and that's it.

What other advice do I have?

Start small and simple. Work with the Vectra support team.

The solution’s ability to reduce false positives and help us focus on the highest-risk threats is the tricky part because we are still doing the filtering. The things it sees are out of the ordinary and anomalous. In our company, we have a lot of anomalous behavior, so it's not the tool. Vectra is doing what it's supposed to do, but we need to figure out whether that anomalous behavior is normal for our company.

The majority of the findings are misconfigurations of servers and applications. That's the majority of things that I'm investigating at the moment. These are not security risks, but need to be addressed. We have more of those than I expected, which is good, but not part of my job. While it's good that Vectra detects misconfiguratons, there are not our primary goal.

The solution is an eight (out of 10).

We don't investigate our cloud at the moment.

Which deployment model are you using for this solution?

On-premises

Read 12 reviews of Vectra AI

See All Reviews