IT Central Station
Unbiased reviews from the tech community

Case Study

Vectra AI

reviewer1439937

Operational Security Manager 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?

Vectra was deployed to give us a view of what is happening on the user network. It helps us to check what is being done by users, if that is compliant with our policies, and if what they're doing is dangerous. It covers cyber security stuff, such as detecting bad proxies, malware infections, and using packet defense on strange behaviors, but it can also be used to help with the assessment of compliance and how my policies will apply.

We also use Vectra to administer servers and for accessing restricted networks.

There are on-prem modules, which are called Cognito Detect, the NDR/IDS solution, which captures traffic. We also have the SaaS data lake, and we also have the Cognito Detect for Office 365, which is a SaaS-type sensor within the O365 cloud.

How has it helped my organization?

If we didn't have Vectra and the Detect for Office 365, it would be very difficult to know if our Office 365 was compromised. We tried, in the past, to do it with a SIEM solution consuming Office 365 logs and it was really time-consuming. The Office 365 Detect solution has the exact same "mindset" as the Detect solution for networks. It's almost like we can deploy it in the fire-and-forget mode. You deploy the solution and everything is configured. You have all the relevant alerts out-of-the-box. If you want to, you could tweak, configure, contextualize, and rewrite the parser, because some things might be out of date, and customize the solution. For a big company with a large team it might be feasible, but for small companies, it's an absolute showstopper. The Detect for Office

365 gives us a lot of visibility and I'm very pleased with the tool.

We use three services from Vectra: Cognito Detect, Detect for Office 365, and Cognito Recall, and we are leveraging all these services within the SOC team to have proper assessments. We even use these tools to prepare the new use cases that we want to implement into our SIEM solution. Recall stores all the metadata that is brought up from Cognito Detect at a central point, data-lake style, with an elastic stack and a Kibana interface available for everybody. Using this, we can try to see what are the general steps. Without this, I would not have been able to have my SOC analyst do the job. Creating a data lake for cyber security would be too expensive and too time-consuming to develop, deploy, and maintain. But with this solution, I have a lot of insight into my network.

An additional thing that is very convenient with the Recall and Detect interfaces is that you can do use cases involving individuals in Recall and have them triggered in Detect. For example, we found ways to track down if users are trying to bypass proxies, which might be quite a mess in a network. We found a type of search within Recall and have it triggering alerts in Detect. As a result, things can be managed.

It's so efficient that I'm thinking about removing my SIEM solution from our organization. Ours is a small organization and having a SIEM solution is really time-consuming. It needs regular attention to properly maintain it, to keep it up and running, consume all the logs, etc. And the value that it's bringing is currently pretty low. If I have to reduce costs, I will cut costs on my SIEM solution, not on Vectra.

The solution also provides visibility into behaviors across the full life cycle of an attack in our network, beyond just the internet gateway. It provides a lot of insight on how an attack might be coming. There are multiple phases of an attack that can be detected. And there is a new feature where it can even consume intelligence feeds from Vectra, and we can also push our own threat-intelligence feeds, although these have to be tested. The behavioral model of the Detect solution also covers major malware and CryptoLockers. I know it's working. We tested some cases and they showed properly in the tool. I'm quite reassured.

It triages threats and correlates them with compromised host devices. One of the convenient things about Detect is that it can be used by almost anybody. It's very clear. It's quite self-explanatory. It shows quadrants that state what is low-risk and what is high-risk. It is able to automatically pinpoint where to look. Every time we have had an internal pen test campaign, the old pen test workstation has popped up right away in the high-risk quadrant, in a matter of seconds. To filter out false positives it can also provide rules that state, "Okay, this is the standard behavior. This subnet or this workstation can do this type of thing." That means we can triage automatically. It also has some features which aren't so obvious, because

they are hidden within the interface, to help you to define triage rules and lower the number of alerts. It looks at all your threat or alert landscapes, and says, "Okay, you have many alerts coming from these types of things, so this group of workstations is using this type of service. Consider defining a new, automated triage rule to reduce the number of alerts."

To give you numbers, with my SIEM I'm monitoring some IDS stuff within my network. Everything is concentrated within my SIEM. From my entire site, IDS is giving me about 5,000 more alerts than my Vectra solution. Of course it will depend on how it is configured and what types of alerts it is meant to detect, but Vectra is humanly manageable. You don't have to add something to make the triage manageable, using some time-consuming fine-tuning of the solution, requiring expertise. This is really a strong point with Vectra. You deploy it, and everything is automatically done and you have very few alerts.

Its ability to reduce false positives and help us focus on the highest-risk threats is quite amazing. I don't know how they made their behavioral or detection models, but they're very efficient. Each alert is scored with a probability and a criticality. Using this combination, it provides you insights on alerts and the risks related to alerts or to workstations. For example, a workstation that has a large number of low-criticality alerts might be pinpointed as a critical workstation to have a look at. In fact, in the previous pen test we launched, the guys were aware that the Vectra solution was deployed so they tried some less obvious tests, by not crawling all the domain controllers, and things like that. Because there were multiple, small alerts, workstations were pinpointed as being in the high-risk quadrant. This capability is honestly quite amazing.

And, of course, it has reduced the security analyst workload in our organization, on the one hand, but on the other it has increased it. It reduces the amount of attention analysts have to pay to things because they rely on the tool to do the job. We have confidence in its capability to detect and warn only on specific things of interest. But it also increases the workload because, as the tool is quite interesting to use, my guys tend to spend some time in Recall to check and fix things and to try to define new use cases. Previously, I had four analysts in my shop, and every one of them was monitoring everything that was happening on the network and in the company on a daily basis. Now, I have one analyst who is specialized in Vectra and who is using it more than the others. He is focusing on tweaking the rules and trying to find new detections. It brings us new opportunities, in fact. But it has really reduced the workload around NDS.

In addition, it has helped move work from our Tier 2 to our Tier 1 analysts. Previously, with my old IDS, all the detection had to be cross-checked multiple times before we knew if it was something really dangerous or if it was a false positive or a misconfiguration. Now, all the

intelligence steps are done by the tool. It does happen that we sometimes see a false positive within the tool, but one well-trained analyst can handle the tool. I would say about 20 to 30 percent of work has moved from our Tier 2 to our Tier 1 analysts, at a global level. If I focus on only the network detections, by changing all my IDS to Vectra, the number is something like more than 90 percent.

It has increased our security efficiency. If I wanted to have the same type of coverage without Vectra, I would need to almost double the size of my team. We are a small company and my team has five guys in our SOC for monitoring and Tier 1 and Tier 2.

It reduces the time it takes for us to respond to attacks. It's quite difficult to say by how much. It depends on the detections and threat types. Previously, we had an antivirus that was warning us about malicious files that were deployed on a workstation within one year. Now, we can detect it within a few minutes, so the response time can be greatly enhanced. And the response time on a high-criticality incident would go from four hours to one hour.

What is most valuable?

The most valuable feature for Cognito Detect, the main solution, is that external IDS's create a lot of alerts. When I say a lot of alerts I really mean a lot of alerts. Vectra, on the other hand, contextualizes everything, reducing the number of alerts and pinpointing only the things of interest. This is a key feature for me. Because of this, a non-trained analyst can use it almost right away.

It's very efficient. It can correlate multiple sources of alerts and process them through specific modules. For example, it has some specific patterns to detect data exfiltration and it can pinpoint, in a single area, which stations have exfiltrated data, have gathered data, and from which server at which time frame and with which account. It indicates which server the data is sent to, which websites, and when. It's very effective at concentrating and consolidating all the information. If, at one point in time, multiple workstations are reaching some specific website and it seems to be suspicious, it can also create detection campaigns with all the linked assets. Within a single alert you can see all the things that are linked to the alert: the domains, the workstation involved, the IPs, the subnets, and whatever information you might need.

The key feature for me for Detect for Office 365 is that it can also concentrate all the information and detection at one point, the same as the network solution does. This is the key feature for me because, while accessing data from Office 365 is possible using Microsoft interfaces, they are not really user-friendly and are quite confusing to use. But Detect for Office 365 is aggregating all the info, and it's only the interesting stuff.

We are still in the process of deploying the features of Detect for Office 365, but currently it

helps us see mailboxes' configurations. For example, the boss of the company had his mailbox reconfigured by an employee who added some other people with the right to send emails on his behalf, and it was a misconfiguration. The solution was able to pinpoint it. Without it, we would never have been able to see that. The eDiscovery can track down all the accesses and it even helped us to open an incident at Microsoft because some discoveries were made by an employee that were not present in the eDiscovery console on the protection portal from Office 365. That was pinpointed by Vectra. After asking the user, he showed that he was doing some stuff without having the proper rights to do so. We were able to mitigate this bit of risk.

It also correlates behaviors in our network and data centers with behaviors we see in our cloud environment. When we first deployed Vectra, I wanted to cross-check the behavioral detection. After cross-checking everything, I saw that everything was quite relevant. On the behavioral side, the Office 365 module can alert us if an employee is trying to authenticate using non-standard authentication methods, such as validating an SMS as a second factor or authenticating on the VPN instead of the standard way. The behavioral model is quite efficient and quite well deployed.

What needs improvement?

Vectra is still limited to packet management. It's only monitoring packet exchanges. While it can see a lot of things, it can't see everything, depending on where it's deployed. It has its limits and that's why I still have my SIEM.

I am in contact with the Vectra team, if not weekly then on a monthly basis, to propose improvements. For the time being, the main improvement I can see would be to integrate with more external solutions. Since Vectra provides an API, that should be quite easy to handle. For example, we're using an open source ticketing system within our team and I want to have it handled properly by Vectra. We'll go forward on that with the API.

Another area for improvement that I have pinpointed is that the Office 365 solution and the Detect solution cannot match the same users. That means we have two "different worlds" currently, the world from Office 365, which is bringing alerts based on users' emails and email addresses. And we have the network world, which is bringing an Active Directory view. On the one hand we are seeing emails or email addresses, and on the other hand we are seeing things like logons on to the domain controller. From time to time, it does not match and the tool cannot currently cross-check this info and consolidate everything. I would like to be able to see that detection related to one workstation and covering a user: what he is using, what services he is using, and what he did with his Office 365 and configuration. That would help.

Another major feature would be to have all logs

pushed to Cognito Detect, and all these logs should be also pushed to Recall. Currently, within Recall, I can't call up the Office 365 detections and I would love to do so.

The last point would be an automated IoT threat feed consumption by the tool.

For how long have I used the solution?

I have been using Vectra for two years.

What do I think about the stability of the solution?

The stability is absolutely flawless. The last time it was rebooted was almost two years ago.

The only thing we have seen was some interruption in log feeding to the Recall instance, the SaaS solution. I had a quick call with a product manager in Europe and he was very keen to share information about this issue and willing to improve it.

So, within two years we have faced one stability incident. This incident lasted less than two hours and it was not on the monitoring solution but more on the data lake solution.

What do I think about the scalability of the solution?

The scalability is very good. From the financial perspective, we are not limited by the number of sensors. We can deploy as many virtual sensors as we want. The key factor is the IP addresses that are being monitored. In terms of technical scalability, we have one brain appliance, one very big sensor, and multiple virtual sensors, and I don't see any limits with this solution.

We are currently using all the things that it's possible to use in this solution. One thing I like with Vectra is that it's updated very frequently. Almost every month new features are popping up: new detections, new dashboards, new ways to handle things. That's quite good. I work with our SOC team so that they can use everything right away.

How are customer service and technical support?

The tech support is surprisingly good. We had questions, we faced some slight issues, and we always got very quick answers. Things are taken into account within a few minutes and answers usually come in less than two hours.

How was the initial setup?

To deploy Recall, which is the data lake in SaaS, or to deploy the Office 365 sensor, it was effortless. It was just a quick call and, within minutes, everything was set up.

It was set up the same way the solution is behaving. It's a turnkey solution. You deploy it and everything works. The configuration steps are minimal. It's exactly the same for the SaaS

solution. You deploy the tool and you just have to accept and do very basic configuration. For Office 365, you have to grant rights for the sensors to be able to consume API logs and so on. You grant the rights and everything is properly set up. It's exactly the same for Recall. It was a matter of minutes, and not a matter of days and painful configurations.

In terms of maintenance it is very easy and takes no time. It's self-maintaining, aside from checking if backups have properly ended. And in terms of deployment, when we add a network segment, we have to work a bit to determine where to deploy the new sensors, but the deployment model is quite easy. The Vectra console is providing the OVA to provide a virtual sensor for deployment. It can also automate the deployment of the sensor if you link it with vCenter, which we have not done. But it's very easy. It's absolutely not time-consuming.

If I compare the deployment time to other solutions, it's way easier and way quicker. If I compare it to my standard IDS, in terms of deployment and coverage, it's twice or three times better.

What about the implementation team?

We were in contact with Vectra a lot at the beginning to plan the deployment, to check if everything was properly set up. But the solution is quite easy to set up. The next decisions we had were focused on how to enhance the solution: what seemed to be missing from the tool and what we needed for better efficiency.

The guys from Vectra were more providing guidance in terms of where the sensors needed to be deployed and that was about it.

We had a third-party integrator, Nomios, that provided the appliances, but they did not do anything aside from the delivery of appliances to our building. Our team took the hardware and racked it into the data center on its own. With just a basic PDF, we set up the tool within minutes. The integrator was quite unnecessary.

Nomios are nice guys, but we have deployed some of other solutions with them and we were not so happy about the extra fees. We were not the only ones who were not happy about that. We tried to deploy the ForeScout products with Nomios and it was quite a mess. But they have helped us with other topics and they have been quite efficient with those. So they are good on some things and on other things they are not good.

What was our ROI?

It's ineffective to speak just about the cost of the solution, because all the solutions are costly. They are too costly if we are only looking at them from a cost perspective. But if I look at the value I can extract from every Euro that I spend on Vectra, and compare it to every Euro I spend on other solutions, the return on investment on Vectra is way better.

ROI is not measurable in my setup, but I can tell

you that Vectra is way more cost-efficient than my other solution. The other solution is not expensive, but it's very time-consuming and the hardware on which it's running it's quite expensive. If I look at the global picture, Vectra is three or four times more cost-efficient than my other solution.

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

The pricing is very good. It's less expensive than many of the tools out there.

Which other solutions did I evaluate?

I evaluated Darktrace but it wasn't so good. Vectra's capabilities in pinpointing things of interest are way better. With Darktrace, it is like they put a skin of Kibana on some standard IDS stuff.

Vectra enables us to answer investigative questions that other solutions are unable to address. It provides an explanation of why it has detected something, every time, and always provides insights about these detections. That's very helpful. Within the tool, you always have small question marks that you click on and you have a whole explanation of everything that has been detected: Why has it been detected and what work is the recommended course of action. This approach is very helpful because I know that if I ask somebody new, within our team, to use Vectra, I don't have to spend months or days in training for him to be able to handle the solution properly. It's guided everywhere. It's very easy to use.

What other advice do I have?

Do not be afraid to link Vectra to the domain controller, because doing so can bring a lot of value. It can provide a lot of information. It gets everything from the domain controller and that is very efficient.

You don't need any specialized skills to deploy or use Vectra. It's very intuitive and it's very efficient.

We are in the process of deploying the solution’s Privileged Account Analytics for detecting issues with privileged accounts. We are using specific accounts to know whether they have reached some servers. It's quite easy with all these tools to check whether or not a given access to a server is a legitimate one or not.

We don't use the Power Automate functionality in our company, but I was very convinced by their demonstration, and an analyst in my team played with it a bit to check whether or not it was working properly. These are mostly advanced cases for companies that are using Office 365 in a mature manner, which is not the case for our company at the moment.

In our company, less than 10 people are

using the Detect solution, and five or six people are using Recall. But we are also extracting reports that are provided to 15 to 20 people.

Read 12 reviews of Vectra AI

See All Reviews