IT Central Station
Unbiased reviews from the tech community

Case Study

Vectra AI

reviewer1302852

Sr. Specialist - Enterprise Security at a mining and metals company with 5,001-10,000 employees

Review by a Real User
Verified by IT Central Station

What is our primary use case?

Our main intention was to see what type of visibility, in terms of detections, Vectra could give us.

We use it on both our manufacturing perimeter and at the internet perimeter. That's where we have placed the devices. We have placed it across four sites, two in UAE and two outside UAE.

How has it helped my organization?

What we have seen over the course of the three to four months it has been in place is that it has not found anything bad. That's good news because nothing specific has happened. But we have identified a lot of misconfigurations as well as some information on how applications are working, which was not known earlier. The misconfigurations that became known because of Vectra have been corrected.

It has given us the opportunity to understand some of the applications better than we had understood them before because some of the detections required triage and, while triaging, or in that investigation, we found how applications work. That is one of the main benefits.

We did a red team penetration exercise and almost all the pen activities were picked up by Vectra. That is another big benefit that we have seen through the deployment of the device.

Apart from the network traffic, a lot of the privileged accounts get monitored. It focuses on the service, the machine, and the account. We have seen many of the privileged accounts flagged with alerts whenever they're doing any

activity which they do not normally do. We can see that it is the admin accounts or our support team accounts where the activity is happening. It is important because any privileged access which sees increased activity becomes a cause for suspicion. It's something that we need to be watchful for. It's a very useful feature because a privileged account can propagate more easily than an account that is not privileged.

These are all examples of the kind of information which is of great value, information that we didn't have earlier.

The detections, as well as the host ratings, allow us to focus in cases where we are pressured for time and need to do something immediately. We can focus on the critical and high hosts, or on the detections that have a very high score. If you do a good job in the rules and policy configuration, the alerts are not too numerous. A person can easily focus on all the alerts. But as of now we focus on the critical, high, and medium. The scoring and the correlation really help in focusing the security operations.While I wouldn't say Vectra AI has reduced our security analyst's workload, it allows him to focus. It's a new tool and it's an additional tool. It's not like we implemented this tool and removed another one. It doesn't necessarily reduce his total time, but what it definitely does is it allows him to prioritize more quickly. Previously, he would be looking at all the other tools that we have. Here, it allows him to focus so things of serious concern can be targeted much faster and earlier. The existing tools remain. But Vectra is something to help give more visibility and focus. In that sense, it saves his time. Vectra is very good for automated threat-hunting, so you get to pick out things faster. All the other tools give you a volume of data and you have to do the threat-hunting manually.Also, the technical expertise required to do the hunting part is much less now, because the tool does it for you. I wouldn't say that it has moved work from tier 2 to tier 1, but both of them can use their time and efforts for resolving problems rather than searching for actual threats. You cannot do away with tier 2 people, but they can have a more focused approach, and the tier 1 people can do less. It reduces the work involved in all their jobs.

In addition, it has definitely increased our security efficiency. The red team exercise is a very clear-cut example of how efficiency has been enhanced, because none of the other tools picked these things up. Vectra was the only tool that did.

It makes our workforce more efficient, and makes them target the actual threats, and prioritizes their efforts and attention. Whether that eventually leads to needing fewer people is a different question. Quantifying it into a manpower piece is probably more an HR issue. But improved efficiency is definitely what it provides. If I needed three or four tier 2 people before, I can manage with one or two now.

And Vectra has definitely reduced the time it takes us to respond to attacks. It's a significant reduction in time. In some cases, the key aspect

is that, more than saving time, it detects things which other tools don't. It helps us find things before they actually cause damage. The other tools are more reactive. If your IPS and your signatures are getting hit, then you're already targeted. What Vectra achieves is that it alerts us at the initial phase, during the pre-damage phase. During the red team exercise we had, it alerted us at their initial recon phase, before they actually did anything. So more than saving time, it helps prevent an attack.

What is most valuable?

The solution's ability to reduce alerts, by rolling up numerous alerts to create a single incident or campaign, helps in that it collapses all the events to a particular host, or a particular detection to a set of hosts. So it doesn't generate too many alerts. By and large, whatever alerts it generates are actionable, and actionable within the day. With the triaging, things are improving more and more because, once we identify and investigate and determine that something is normal, or that it is a misconfiguration and we correct it, in either of these two instances, gradually the number of alerts is dropping. Recently, some new features have been introduced in the newer versions, like the Kerberos ticketing feature. That, obviously, has led to an initial spike in the number of tickets because that feature was not there. It was introduced less than a month back. Otherwise, the tickets have been decreasing, and almost all the tickets that it generates need investigation. It has very rarely been the situation that a ticket has been raised and we found that it was not unique information.

Also, we have seen a lot of detections that are not related to the network. Where we have gained extra value in terms of the internet is during data exfiltration and suspicious domains access.

The detections focus on the host, and the host's score is dependent on how many detections it triggers. We have seen with many of our probing tools, without triaging, that these hosts pretty quickly come into the high-threat quadrant. Its intelligence comes from identifying vulnerable hosts along with the triaging part. That's something that we have seen.

What needs improvement?

One thing which I have found where there could be improvement is with regard to the architecture, a little bit: how the brains and sensors function. It needs more flexibility with regard to the brain. If there were some flexibility in that regard, that would be helpful, because changing the mode of the brain is complex. In some cases, the change is permanent. You cannot revert it. I would like to see greater flexibility in doing HA without having to buy more boxes just to do it.

Another area they could, perhaps, look at is with OT (operational technology) specifically. Vectra is very specific to IT-related threats. It really doesn't have OT in its focus. We are using

another tool for that, but maybe that is another area they can consider venturing into.

It's being used by my team of four or five people. Once we hand it over to operations, then the team size will increase significantly. It will grow to about 10 to 15 people.

For how long have I used the solution?

We have been using Vectra AI four about four months.

What do I think about the stability of the solution?

Stability-wise, we've not had any issues, although it has only been three or four months. We had some slight bugs in there, bugs that were related to the triaging and how we used the conditions. But stability-wise, we've had no problem.

There were some software issues, bugs, but then nothing major. There were minor cosmetic and syntax-based issues while raising the conditions. Apart from that, no issues with the stability.

What do I think about the scalability of the solution?

Currently we are in the process of expanding it to two more remote sites. One is in West Africa, in Guinea, and another one in the U.S. Those are more recent deployments, in place less than a month. We are in the process of creating the policies, and triaging, and investigations for those. That's ongoing. With those sites, the benefit realization is still pending because we just started the traffic loading.

The scalability part is where the architecture comes in. That's one of the areas for improvement that I would like to recommend. Unless you have dedicated brains doing anything other than brain functions, it doesn't become scalable. If you have a brain in mixed mode, your scalability is limited. Also, the brain's capacity gets reduced based on its function, so if it's in mixed mode, the capacity is less. If it's in brain mode, the capacity is more. If it's in sensor mode, the capacity is different. It makes scalability difficult. Unless you go for two big brains with your highest capacity device and then you keep adding.

When I spoke to our internal success team at Vectra, they mentioned that this is something that they're planning to fix in the near future with an upgrade.

How are customer service and technical support?

Whenever we have raised issues we have gotten timely responses. Getting support is fairly easy compared to some of the other technologies that we have. A simple email is sufficient to get attention from their support

team. They have a remote access feature wherein we don't necessarily have to give a WebEx. We just simply enable the remote access on the device, and the remote team can log in, and have a look, and understand what the problem is.

How was the initial setup?

The problem was the architecture. Once we arrived at an architecture, it was simple. What takes time is to build the architecture plan because of the way the brains work. We had to agree on a design. Once you agree on the architecture, the implementation is pretty straightforward.

The initial architecture design took some time, a week or so. The implementation was done within a day.

Our implementation strategy was to have an HA setup for each site. We put two brains into mixed mode, but then we found out that if we put it in mixed mode, HA is not possible. So we set it up as a standby and we configured manual scripts to transfer the file from one brain to the other brain. That's how we are managing it now. If we want to go live on the standby brain, we just import the configuration and go live, if there is a failure.

It's a little bit manual process for us. If it has to be automated, I believe the brains cannot be in mixed mode. That was where we faced the initial problem, I mean, for the architecture part. So we have two brains configured in mixed mode and we have a couple of sensors on the OT side, sensors that are talking to these brains. The sensors are there in the OT connectivity, the active or standby firewalls, and this is repeated on the other site as well.

Two or three people are enough for the deployment. They should have a sound understanding of the network and an idea of how the architecture and the applications function. One person from the architecture team and one person from the network or security team are sufficient to understand how to get maximum utilization from Vectra.

What was our ROI?

In terms of visibility and security improvement, we have definitely seen a return on our investment.

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

We have a one-year subscription that covers support and everything. There is no other overhead.

Which other solutions did I evaluate?

We evaluated Darktrace, in addition to Vectra, each in a PoC. We chose Vectra because the

things that Vectra picked up were far more useful, and necessary from an enterprise point of view. Darktrace was a bit noisier.

What other advice do I have?

One thing we have learned using Vectra is that anomaly detection is a critical component of security; a non-signature-based technology is very critical. It helps pick up things that other tools, which are more focused on active threats, will miss. That is one major lesson that we have picked up from Vectra.

My advice would be that you need to focus, because the licensing is based heavily on IPs and area of coverage, although predominantly IPs. You need to have a very clear idea of what areas you want to cover, and plan according to that. Full coverage, sometimes, may not be practical because, since it's a detection tool, covering everything for large organizations is complicated. Focus on critical areas first, and then expand later on.

Also, the architecture part needs to be discussed and finalized early on, because there is a limited flexibility, depending on which model you choose to take.

The solution captures network metadata at scale and enriches it with security information, but the full realization of that will come with Cognito Stream, which we have yet to implement. Right now we are on Cognito Detect. Cognito Stream is something that we are working on implementing, hopefully within the next month or so. Once that comes online, the enriched metadata will have greater value. As of now, the value is there and it's inside Vectra, but we don't see that information — such as Kerberos tokens, or certificates, or what the encryption is — unless it leads to a detection. Only in that event do we currently see that information.

The Cognito Stream can feed into our SIEM and then we will have rich information about all the metadata which Vectra has in our data lake.

Which deployment model are you using for this solution?

On-premises

Read 12 reviews of Vectra AI

See All Reviews