IT Central Station
Unbiased reviews from the tech community

Case Study

Vectra AI

VECTRA

reviewer1259193

Head of Information Security at a insurance company with 1,001-5,000 employees

✓ Review by a Real User
✓ Verified by IT Central Station

What is our primary use case?

One of the biggest things is the visibility of stopping or identifying any infection as soon as possible. In this case, if someone downloads something malicious to their workstation, we have a number of controls in place. However, it wasn't so much the endpoint. It was the spreading of a worm type scenario or a WannaCry type thing. Anything that could potentially spread after the initial infection, which is where we wanted to come in and get that visibility.

It was key for us to have something that we could use for identifying as soon as possible, which would be call center initiated. That was probably our biggest thing: To push it in that direction, as we're a regulated company from the FCA. They drive us continually for improvement and behavioral analysis. Network analysis sort of falls into that bucket.

We already have a SIEM, which some people would argue gives us a lot of that visibility. It doesn't tend to give it the focus that we need. From Vectra, we get a lot of alerts of, "This is happening," or, "This is unusual." This is a lot easier than waiting for a couple of logs to come in, then a bit of AI logic at the back of it to potentially push it in that direction. It's very much for us to get a view of a potential attack, then deal with it as quickly as possible. To pinpoint where it's coming from, and where it is going to go.

One of the biggest things that I wanted to ensure is that it covered our call centers because that is where I see my biggest risk. So, I was really key on getting sensors across all geographic locations within the UK and in all of our small communication rooms.

It is all on-premise. We have a number of call centers spread around the UK. We look at all east-west traffic, as well as north-south. It all goes into our brain in our data center. We do have some branches out in Azure, but we're waiting on the new plugin that they are trying to develop. We are just starting in on our cloud journey and most of our infrastructure is in still private cloud. We haven't really gotten to the point where we have public cloud.

We're up-to-date, but I don't know the exact version number that we are on.

How has it helped my organization?

The key improvement for us were:

The additional monitoring 24/7, and using the high fidelity alerting from Vectra rather than SIEM, This was our biggest change. We have managed to leverage that rather more than our SIEM, which just throws out loads of spam.   The FCA requirements to build on behavior monitoring. The use case of the call center with its high turnaround of staff who are perhaps not as clued in or engaged in our user awareness program as they could be. Lack of end user deployment is another big improvement. We wanted something that was easy to deploy, or get up and running really quickly. It took a couple of weeks to rid of the alerts that we didn't want, but the actual involvement from the network teams was minimal, which was really good for us because we just don't have the resource to spend a lot of time trying to configure devices. We use the solution’s Privileged Account Analytics for detecting issues with privileged accounts. Although I haven't seen a huge amount of alerts. We have a quarterly QBR, and they mentioned it the day before the QBR and noticed an alert pop up.

One of the key things for us is we have an annual pen test (an internal one), that's not as involved as a full red team. But, it's enough for the pen test to sit with the SOC guys, then we put the different tool sets together, what they're doing, and how that reacts to our Vectra, SIEM, and endpoint AV. To see what picks up where, so it gives us an ability to check those tool sets that we have.

From a Vectra point, it will pick up a number of different things. But, it will also miss a number of different things. That's how pen testers work. They work covertly. So, it's really good for us to see what we can do and what we can't. Then, that feedback goes back to Vectra. We say, "Okay, well why didn't we pick up this?" They'll come up with a reason or they'll take it away and find something out about it. That's really good and a nice part of the service. We get to check to make sure the tool sets are working, but we also provide feedback and they're very open to that type of feedback.

I believe the solution has increased our security efficiency. It's hard to prove without having a direct attack. But, I get challenged about ransomware from my board, to say, "How do we defend against ransomware?" That's a big topic.

One of the key things was when Vectra went in, it saw a developer run a script, which essentially changed the names for a number of files and put a different extension on, but they were doing some development type work. That's how their script ran, and it identified that as ransomware, which is a great thing to say.

Although there was no encryption or malice involved, it did create new files, rename files, and delete old files, which essentially is what ransomware does anyway. It followed the same sort of logic to it, I can report that back. "We do have some protection. It wouldn't stop it. But we could limit the amount of damage that it may do."

I don't know about other companies, but I get the feeling most people look to identify rather than block. We're not a high-end bank. We are not going to stop people working. We're going to investigate what they're trying to do. That's just our risk appetite. We have to work. Unless it's absolutely 100 percent, we won't stop them. We would just look at it afterwards. So, all our alerting, we don't have any orchestration at the back of it to say, "Okay, if this happens, then I'm going to play that port in a firewall or I'm going to drop that from there." We won't do that. Humans will all be part of that process. We'll get a call, then we will make a crisis management team decision, etc. That's how we operate.

If, for instance, our AV doesn't pick it up. I think that is where Vectra will come in. So, if somebody gets infected and maybe hasn't picked it up. That's where, if that worm spread and our endpoint signatures weren't up-to-date, they went into zero day, and nobody knew about it. Vectra would give us that opportunity. It would potentially give us something that would say, "Well, this is not normal. This machine does not communicate with all these other machines like it is now." That's where we see it coming in. It gives us that extra chance to stop a disaster before it happens, or at least limit the amount of potential output of damage that that an incident can do.

Zero days are always very difficult. If the AV vendor doesn't know about it, it's not going to be able to tell me about it, stop it, quarantine it, or do anything. Having a tool set like this, which monitors network traffic for anomalies, it gives us that chance. I can't say that it definitely will pick it up, but there's another opportunity for us to reduce the amount of damage that can be done.

What is most valuable?

It gives us the point of where something is happening, which is the key thing for us. (I know that there is a back-end recall, which probably gives a lot more data, but we don't use that.) We then leverage our SIEM product to provide us logs from those specific sources that it's talking about, giving us that information. It is the accuracy of: It is happening here and on this particular host, then it's going to here to this particular host. It's that focus which is probably the most advantageous to us.

The logic behind Vectra's ability to reduce alerts by rolling up numerous alerts to create a single incident or campaign for investigation grows with severity, as there are additional alerts around that particular host. This is a useful feature rather than spamming alerts. But, we've never really had an issue with a lot of alerts. We really do triage our alerts quite well and have a good understanding of what does what.

One of the key advantages for us is we define a 24/7 service around it. We use far more of Vectra alerts than we do with our SIEM product because we understand that when we get an alert from Vectra we actually need to do something about it. You can't really say you don't get false positives, as the action has happened. It's whether we consider that action as a concern rather than a SIEM that sort of gives you a bit of an idea of, "That may be something you're interested in." Whereas, Vectra says, "This has happened. Is this something you would consider normal?" I think that's the bit that we like. It just says, "Is this normal behavior or isn't this normal?" Then, it's up to us to define whether that is or isn't, which we like.

The solution provides visibility into behaviors across the full lifecycle of an attack in our network, beyond just the internet gateway, because we do east-west traffic. So, it looks at the entire chain across there. We're fortunate enough not to be in a position that we've seen a meaningful attack. When we do have pen testers come in, we can see quite clearly how they pick traffic up and how it develops from a small or medium alert to go to higher severity, then how it adds all those events together to give more visibility.

The solution does a reasonable job of prioritizing threats and correlating them with compromised host devices. We use that as how we react to it, so we leverage their rating system. We are reasonably comfortable with it. At the end of the day, we actually spend a lot of time and effort to tweak it. It's never going to be right for every company because it depends on what your priorities are within the company, but we do leverage what they provide. If it is a high, we will treat it as a high, and we will have SLAs around that. If it's a low, we'll be less concerned, and the events that come out pretty much lead to that. The events that we see and the type of activity going on, it makes sense why it's a low, medium, or high. Just because a techie has done a port scan, that doesn't mean we need to run around shouting, "Who has done this?"

When we originally put it in, it was really quite interesting to see. Picking up the activities from the admin user and what they were doing, then going, "By the way, why have you done that?" Then looking at a scan and going, "Well, how did you know that?" So, it's sort of cool to pick up that type of stuff. We tend to trust what it tells us.

What needs improvement?

Room for improvement depends on how their strategy and roadmap develops, as they have a lot of third-parties that they integrate with, e.g., more orchestration around what alerts and what to do with afterwards. They don't pretend to be working in that space. That is a third-party type activity.

There are always the little things that they could do a bit better, like grouping or triage filters. Clearly, they've taken that onboard and developed those over the course of the last 18 months to two years to put these additional functions in. My guys are constantly saying, "Oh, it'd be useful to do this and useful to do that."

The solution has not reduced the security analyst workload in our organization because we still need to SIEM. Unfortunately, while Vectra, for us, is a brilliant tool for network investigations, giving wonderful visibility, it doesn't go the whole way to replace our SIEM that is needed for compliance. So, I still have the same amount of alerting and logging that I did before. It gives us more defined ability to see incidents, but it doesn't give us enough information to satisfy a PCI or 27001 audit.

For how long have I used the solution?

I have been using this solution for about two years.

What do I think about the stability of the solution?

Interestingly enough, when we first got Vectra, we had a number of problems with it. The guys were all over the solution trying to fix it. It turned out to be a hardware issue. I think they ended up changing their supplier. They just ripped everything out and put a load of new equipment in. This was identified about three months after it being here.

These things happen. There's not a lot you can do about it. However, they were really good and didn't make any excuses, apart from, "It can only be the hardware," which it was. Once they put the new hardware in, everything went really well.

Very few people are required for maintenance. We just generally run the alerts now. I have a guy spend probably less than an hour a day, maybe less than that, putting out fires and alerts. Then we investigate that, depending on its severity. The actual hardware maintenance is nothing. We'll just keep an eye on it or get an alert if an interface has gone down.

What do I think about the scalability of the solution?

One of the biggest things that we wanted to implement was something that was easy to do. Our problem, as well as I'm sure a number of other companies, is the amount of resources to install these new technologies, then how the resource center operates and uses these technologies. It's great having all these additional add-ons here, there and everywhere, but my team is quite small. So, it had to be quite easy. It has to be quite focused. Hence, we went with Vectra.

At the moment, we have a hardware brain and are not near the limit of that. To go from that, I think Vectra was looking at some sort of applied solution, but it would then be a change. So, we're down to limitations of the hardware. I always say, "If we bought a massive company, we would probably have to redesign and architect the solution." At the moment, they made sure that we have some growing room.

Our purchase was a one time thing for the entire company, otherwise we would be leaving ourselves exposed. Just this week, I took a Vectra device up to a new company that we purchased and stuck it in there. It is really that simple. We'll probably end up with a bit of traffic because we will see a lot of new servers and workstations that we have to do triage around.

We have probably 3,500 to 4,000 users across the UK. My team is quite small. I have a couple of guys who are cyber-related.

How are customer service and technical support?

The technical support is brilliant and really responsive. That is probably down to the fact that they are a small company. Their guys respond instantly, normally within the day that they have somebody online and having a look at it, or they're putting it away and the communication is excellent. They will say, "Okay, we'll put it back to the developers," and then they give us updates, which is really efficient.

Vectra is growing at the moment. They support us very well. They do seem to rely on key people. Would my service be the same if they got rid of our technical manager? I don't know. They are a small, close family team, which is really good. Whether that would change when it's a few key people left, I don't know. But I know they are growing as a company as well, so let's hope they scale it in proportion to their customer base. Only time will tell. Other companies I've got at the moment grew too quickly in the services and service suffered as a part of that.

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

It isn't a tool set to replace a current tool set. It's just an additional feature. For me, it has only increased our workload, but that's because we had nothing there before.

We did not previously have a network monitoring solution. We have a toolset that does event log monitoring, but nothing across the network itself. I think we have basic flow visiblity, and the network team use that. However, there is no real way of investigating individual network packets, then using them for anything in particular.

How was the initial setup?

The initial setup was easy. We have multiple sites, so we had to go around and travel to different sites. However, the actual brain was conifgured in a few hours. Once it was up, it was up. The network guys did nothing after that point. My guys probably spent a couple of days, over the course of a month just tweaking it. Then, it gradually goes down as we get a new server pop up, which might add a bit of additional alerting. Once we get a handle on that, then it comes down to something really quite manageable.

The priority for us was to get the main call center up and running at the start. We needed the brain up and do the implementation to see the east-west traffic in our call center. Then, we brought on additional sites, depending on the size around the UK, as we monitored it.

What about the implementation team?

We used the Vectra guys for the implementation. Our technical engineer came in, going into the data centers with our network engineer (or remotely), then set it up.

For the actual deployment within the data center and around the sites, just two people were needed one form each company. After that, it was the configuration of the alerting which took one of the SOC guys suing Vectra for reference.

They provide us a health check and provide us with recommendations on what we need to do every quarter, which is perfect. There is nobody else who does that. That is probably part of the advantage of being a smaller company.

Once every quarter, they'll put health and safety in, and say, "Alright, these are the new functions. This is what you need to turn on. That's not quite working. Those haven't fired. You might want to look at removing those." This is really good to see, because I get a lot of vendors, who once they've sold you a technology, they don't really care. They go, "Yep, there you go." They don't look at what you installed, how you've installed it, provide any recommendations, or look at how it's performing.

This provides me the assurance my SOC guys are doing a good job, we are on top of any changes and the assurance we are getting the most of the solution.

Vectra has pretty much forced this upon us, which is really good because everyone is very busy. Before you know it, the months turn to years and disappear.

What was our ROI?

ROI is a difficult one for security tools. You can argue that if you don't see anything where you did investment, this is the reason to have good security tools: not to have an incident. You only really know when bad things happen, and you're in the middle of it. Otherwise, it's doing what it needs to do to stop or identify an issue in the first place.

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

We are running at about 90,000 pounds per year. The solution is a licensed cost. The hardware that they gave us was pretty much next to nothing. It is the license that we're paying for. I think if we outgrow our current hardware, then we will have a look at bigger hardware or some sort of distribution. I'm sure they have a number of different options for larger companies. I don't see that being a major issue for us in the next three to five years.

We don't have complete visibility because we don't have all of that metadata surrounding it. Sometimes there might be more metadata before, it might be something afterwards, or there might be something missing, but we accept that because we don't have the funds to pay for the additional functionality that it can provide its a trade off.

Which other solutions did I evaluate?

When we started off, apart from money, we had to look at behavioral analysis. We weren't sure where we wanted to go with the solution, whether we wanted to look at the endpoint or network. So, after a RFI, to define which direction we wanted to go, we thought that we would go down the network analysis route. Because we have call centers, there is normally a high turnover of staff. The jobs themselves are quite intense and people move around quite a lot, it was key for us to get some visibility in what those guys are doing. We thought, "Although we do a lot of user awareness and logging, this is probably where our weakest link is." It was a case of somebody potentially clicking on a malicious link, some sort of phishing attack which was probably, or is probably, going to cause us the most pain.We looked at Darktrace and there was another option that dropped out. So, we looked at the main players in that area. We decided on the behavior analysis for network, then we took the top three: Vectra, Darktrace, and another solution.

It came down to Darktrace and Vectra. Darktrace looked much prettier than Vectra, unfortunately the support that we'd heard about and reviews that we read, led to, "Here's the new tool set. Off you go". This is what we didn't want. We wanted somebody to hold our hand, then give us the support we needed to ensure we get the best out of the tool set.

It obviously comes down to price as well and we feel we picked the best product that fitted us. We did quite a lot of due diligence on both. I went to different places that got both installed and got references from both. I firmly believe that both products would have done the job well. However, the support from Vectra along with their customers' references to say how good it was, I think we made made the right decision.

What other advice do I have?

People do a lot more than we actually see. Looking at the test and development guys, sometimes they do things that they don't understand. So, they will do it because it works. The actual things that are behind the scenes are the sort of things that happen, and they don't really understand. If there's something that's really complicated, they're people that have initiated it that don't really know what it is. That is always a problem, because in our sort of company, we have a lot of developers who are doing a lot of coding and things like that, but they're not 100 percent on all the other things that they affect, such as the supporting applications underneath it.

They are making a change on one particular app, but it's using the other apps underneath it to develop that and push that across to something else. All these extra, different steps that they are completely oblivious to where we go, "Actually, you've just done this." They go, "Well, I don't know, I just ran the script over here. I don't know why that would happen." But, it'll do a LDAP lookup or connect to a share. Those are the sort of things that you get a lot of visibility from people who don't understand. So, that can become tricky. That's pretty much par for the course for a lot of security tool sets. Where you have a couple of people who know one particular aspect, but don't really understand everything that's going on. To be fair, IT is a big area. You can't expect everyone to know everything of everything, not when you're not working in a massive IT structure, and the security team is a small department.

You need to be quite key on your business case and what you're expecting from it. Be 100 percent sure on your use cases. It's an excellent tool. It doesn't create a huge amount of overhead, but it is a tool that you need to keep on top of. The more you keep on top of it and get it right at the start, the easier it will make your life going forward. Don't just stick it in, then leave it to whirl away as a lot of people do. You have to spend that bit of extra time, and it's not huge amount of time, and leverage other teams.

The way they do their customer success is really good. There's nothing bad that I've got to say apart from the costs, but nothing's free, is it?

It has to be up there with my favorite security tool set at the moment. I am quite lean on scores, but the solution is definitely nine (out of 10). If I look at all my other security tool sets, this is the one that my guys value the most.

Which deployment model are you using for this solution?

On-premises

Read 12 reviews of Vectra AI

See All Reviews