IT Central Station
Unbiased reviews from the tech community

Case Study

Vectra AI

reviewer1296420

Global Security Operations Manager at a manufacturing company with 5,001-10,000 employees

Review by a Real User
Verified by IT Central Station

What is our primary use case?

We use Vectra with the assumption that our other defensive controls are not working. We rely on it to be able to detect anomalous activities on our network and trigger investigation activities. It's a line of detection assuming that a breach occurred or has been successful in some way. That's our primary use case.

We have it in some of other use cases, like anomalous network activity and detection for things. E.g., we are trying to refine or improve suspicious internal behaviours because we are a development technology company. We have developers doing suspicious things all the time. Therefore, we use it to help us identify when they are not behaving correctly and improve our best practices.

We have it predominantly on-prem, which is a combination of physical and virtual sensors. We also have a very minor element on the cloud where we are trialing a couple of components that are not fully deployed. For the cloud deployment, we are using Azure.

We are on the latest version of Cognito.

How has it helped my organization?

We have a limited use of Vectra Privileged Account Analytics for detecting issues with privileged accounts at the moment. That is primarily due to the fact that our identity management solution is going through a process of improving our privileged account management process, so we are getting a lot of false positives in that area. Once our privilege account management infrastructure is fully in

© 2021 IT Central Station | www.itcentralstation.com

place and live, then we will be taking on more privileged account detections and live SOC detections to investigate. However, at the moment, it has limited applicability.

We have a lot of technically capable people with privilege who are able to do things they should or should not be able to do, as they're not subject-matter experts when it comes to things like security. They may make a decision to implement or download a piece of software, implement a script, or do something that gets the job done for them. However, this opens us up to major security risk. These are the types of activities that the tool has been able to identify, enabling us to improve communication with those individuals or teams so they improve their business process to a more secure or best practice approach. This is a good example of how the solution has enabled us to identify when people are engaging in legitimate risky activities, and we're able to identify and engage with them to reduce risk within the network.

It has enabled our security analysts to have more time to look at other tools. We have many tools in place, and Vectra is just one of them. Their priority will always be to deal with intrusion attempt type of alerts, such as malware compromise or misuse of credentials. Vectra was able to simplify the process of starting a threat hunting or investigation activity on an anomaly. Previously, we weren't able to do this because the amount of alerts and volume of data were just too large. Within our security operations, they can now review large volumes of data that provide us with indicators of compromise or anomalous behaviour.

By reducing false positives, we are able to take on more procedures and processes. We have about seven different tools providing alerts and reporting to the SOC at any one time. These range from network-based to host-based to internet-based alerts and detections. We are more capable to cover the whole spectrum of our tooling. Previously, we were only able to deal with a smaller subset due to the sheer workload.

In some regards, I find that Vectra probably create more investigative questions. E.g., we need to find answers from other solutions. So, it is raising more questions than it is specifically answering. However, without Vectra, we wouldn't know the questions to ask in the first place. We wouldn't know what anomalies were occurring on our network.

Vectra data provides us with an element of enrichment for other detections. For example, if we see a detection going onto a single host, we could then look at that activity in Vectra to see whether there are suspicious detections occurring. This would give us the high percentage of confidence that the compromise was more severe than a normal malware alert, e.g., destructive malware or commander control malware enabling someone to pivot horizontally across the network. Vectra provides us with that insight. This enables us to build up an enriched view quickly.

© 2021 IT Central Station | www.itcentralstation.com

What is most valuable?

One of the most valuable features of the platform is its ability to provide you with aggregated risk scores based on impact and certainty of threats being detected. This is both applied to individual and host detections. This is important because it enables us to use this platform to prioritize the most likely imminent threats. So, it reduces alert fatigue follow ups for security operation center analysts. It also provides us with an ability to prioritize limited resources.

It aggregates information on a host and host basis so you can look at individual detections and how they are occurring over time. Then, you can have a look at the host scores too. One of the useful elements of that is it is able to aggregate scores together to give you a realistic view of the current risk that the host plays in your network. It also ages out detections over time. Then, if that host is not been seeing doing anything else that fits into suspicious detection, it will reduce its risk score and fall off of the quadrant where you are monitoring critical content for hosts that you're trying to detect.

When you are analyzing and triaging detections and looking for detection patterns, you are able to create filters and triage detections out. Then, in the future, those types of business usual or expected network behaviours don't create false positive triggers which would then impact risk scores.

Without the detection activities that come from Vectra, we wouldn't have been able to identify the true cause of an event's severity by relying on other tools. This would have slipped under the radar or taken a dedicated analyst days to look for it. Whereas, Vectra can aggregate the risk of multiple detections, and we are able to identify and find them within a couple of hours.

What needs improvement?

You are always limited with visibility on the host due to the fact that it is a network based tool. It gives you visibility on certain elements of the attack path, but it doesn't necessarily give you visibility on everything. Specifically, the initial intrusion side of things that doesn't necessarily see the initial compromise. It doesn't see stuff that goes on the host, such as where scripts are run. Even though you are seeing traffic, it doesn't necessarily see the malicious payload. Therefore, it's very difficult for it to identify these type of host-driven complex attacks.

It only shows us a view of suspicious behaviours. It doesn't show us a view of key or regularly attacked company targets. This could be because we don't have one of the other tools or products that Vectra provides, such as Stream or Recall.

My challenge with the detection alerting platform, Cognito, is it tells us this host is behaving suspiciously and is targeting these other machines, but it won't give you a view when a host is the target of multiple attacks. This because you may have a key

© 2021 IT Central Station | www.itcentralstation.com

assets, such as domain controllers or configuration management servers. These are key assets which may get targeted. If you're a savvy attacker, you spread out your attack across multiple sources to try and hide them across the network. That is where the solution falls a bit short. It is trying to build that chain of relationships across detections and also trying to show detections from a perspective of a victim rather than the perspective of an attacker. I have expressed these concerns to Vectra and they are currently in as feature requests.

There is another feature in place which takes additional data feeds, such as DHCP IP allocation data. Their inputs are taken from Windows event logs, and that's the format they have in place. They use that to provide them with a more accurate view of host identities. If you are only relying on IP addresses, and IP addresses change over time, it's sometimes very difficult to show a consistent view of a system behaviour over time, as the IP can change per month. Unfortunately, because their DHCP data is taken from Windows host events and our DHCP data is taken from a Palo Alto system that generates the IP leasing, the formats are incompatible. I think taking different formats for that type of data is something else we have a feature request in for. At the moment, we don't have an accurate view, or confidence, that they are resolving when an IP address changes from host to host. So, we may be missing an accurate view of risk on some of those hosts.

We also have the same problem with VPN and Citrix. E.g., if you're on the network and on IP address A, then you come in via the VPN, you're now on IP address B. Thus, if you're spreading your suspicious behaviour across both the internal network and VPN, then across Citrix, we don't get to join all that information up. They are seen as three different systems, so it causes a bit of a problem trying to correlate that type of event data.

For how long have I used the solution?

If you include the proof of concept, I have been using Vectra for three years.

What do I think about the stability of the solution?

There are no concerns regarding the stability. It seems to be very reliable. I've had one sensor in two and a half years become corrupt and need to be rebuilt. That's it.

Day-to-day maintenance takes half an FTE to one FTE a day. There is no maintenance really required on the platform. All we need to do is monitor for when a health alarm occurs (a sensor is not working), then we raise the relevant request with the teams to investigate. Maintaining the health of the platform requires a feed into our operations team to be able to look at our monitor to determine when the health is degrading. Doing general health, like detection filters, triage filters, reviewing, looking for

© 2021 IT Central Station | www.itcentralstation.com

patterns and anomalies, and creating new filters, needs a daily dedicated FTE.

What do I think about the scalability of the solution?

The scalability is brilliant. It is able to cope with virtual sensors. You can increase the hardware that supports the image and it will work with the high bandwidth of the data going through. There are no concerns in terms of the scalability.

It does create capture network data at scale because we have it deployed at over a 100 geographically split sites. We have over 8000 users on cloud. So, it's able to deal with the network traffic very easily, providing us with additional information. If we were just relying on things like firewalls and packet capture applications, we wouldn't get to that enrichment of a security context put on top of normal network traffic.

Mainly, there are five people dedicated to using the platform: Tier 2 security analysts and an operations director. However, that is widen out to whomever we are raising the support requirements to, like the Tier 3s. When raised, we also enable the shared link so they can go into the platform and look at the data associated with the detection on that host. So, there is a wider volume of people who use the solution to get information for specifically requested cases.

How are customer service and technical support?

The technical support is very good. They always respond within a short amount of time to provide expert information and have always been helpful in trying to work through problems to find a good solution.

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

Previously, we had a general sensor solution taking logs. We didn't have an equivalent detection platform for our network nor did we have a tool capable of providing us with competent intrusion detection capabilities post-breach. Our main SIEM logging platform was generating over a 1000 alerts a day. It was bloated and unusable when trying to identify events/anomalies that were occurring. Once we implemented Vectra, it was able to give us a refined view and tell us which things we need to prioritize so we were able to reduce our workload from a 1000 alerts a day down to 10.

How was the initial setup?

The initial setup was relatively straightforward. It was pretty much plug and play.

The initial pilot deployment took weeks, but that was because the scope kept on changing. However, the initial deployment only took hours.

© 2021 IT Central Station | www.itcentralstation.com

It has not helped us move work from our Tier 2 to Tier 1 analysts, but this is a fault in our implementation. The structure of our organization hasn't necessarily changed. We don't have Tier 1 security analysts. Therefore, we don't have the capacity or capability for them to deal with these types of detections. We have to leave our Vectra detection and activities with our Tier 2s.

We now have an implementation strategy. We have virtualized sensors in most locations rather than physical sensors. We only have physical sensors in the areas where there is high bandwidth traffic, such as key internal data centers. The virtual centers for local offices are sufficient for the volume of traffic there. We only deploy in areas that are key risks. We also only deploy and monitor network zones which are of significant risk, so we don't monitor our guest WiFi subnet nor do we monitor our development network subnets. Therefore, we keep our segregated networks and zoning structure consistent so we are able to only monitor for priority areas.

What about the implementation team?

Vectra had an engineer come down. They plugged the device in and set it up. Since the firewall rules were already in place, it was working.

Assuming the firewall rules are already in place for the physical sensor, it needs one person plugging it in and putting it into a rack. If it is a virtual sensor, then it is just somebody who can deploy the virtual image onto the virtual infrastructure and switch it on. It takes two dedicated people to deploy. If you have a network team and a server team, then you will need one of each of those skill sets to be able to deploy the tool. It all depends on how your organization is structured.

What was our ROI?

It has increased our security efficiency because we can now do more with the tool. E.g., if we had a data analyst who was creating models and searching the data to identify the same types of the numbers/behaviours within Vectra, we would need at least two or three FTEs.

Vectra has reduced the time it takes us to respond to attacks. In 2019, we conducted a red team activity. The Vectra appliance was able to alert the red team on activity within three hours of the test starting. Prior tests to that, in real life or red team scenarios, we were potentially looking at days. However, we also tightened controls prior to that testing period. While Vectra has done an amazing job in reducing the time to respond, there are so many other things that we also have put in place which have contributed towards it.

Vectra has saved us weeks, if not months, in terms of the ability to identify a breach. Our process has been reduced down to hours, which is a potentially massive return on investment, if

© 2021 IT Central Station | www.itcentralstation.com

we were compromised. From an insurance perspective, the return investment is fantastic.

From an FTE perspective, while it reduces the number of events that we have to look up and the number of alerts, we now have very specific things where we need to ask questions. Therefore, it's creating more work which we weren't capable of doing.

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

At the time of purchase, we found the pricing acceptable. We had an urgency to get something in place because we had a minor breach that occurred at the tail end of 2016 to the beginning of 2017. This indicated we had a lack of ability to detect things on the network. Hence, why we moved quickly to get into the tool in place. We found things like Bitcoin mining and botnets which we closed quickly. In that regard, it was worth the money. Three years later, the license is now due for renewal so we will need to review it and see how competitive it is versus other solutions.

When we implemented the physical sensors, there were costs for support in terms of detection review sessions. We had a monthly session where an analyst would talk through the content, types of detections that they were seeing, etc.

We have a desire to increase our use. However, it all comes down to budget. It's a very expensive tool that is very difficult to prove business support for. We would like to have two separate networks. We have our corporate network and PCI network, which is segregated due to payment processing. We don't have it for deployed in the PCI network. It would be good to have it fully deployed there to provide us with additional monitoring and control, but the cost associated with their licensing model makes it prohibitively expensive to deploy.

Which other solutions did I evaluate?

We did review the marketplace and look around. For example, we looked online at Darktrace, but we didn't run a side by side comparison to see which one would work better.

Vectra was the only tool in which we did a physical pilot or proof of concept. Vectra stood out for its simplicity and the general confidence that I had with the people whom I was engaging and having conversations with at that time. I am very much a people person. If I talk to people and don't get the impression they know what they're talking about, then that will reduce my confidence in their product. E.g., our initial engagement with Darktrace wasn't good enough to provide confidence in their platform, and we had to move quickly.

© 2021 IT Central Station | www.itcentralstation.com

What other advice do I have?

Make sure you have a dedicated resource committed to daily use of the tool. Because the selling point is it frees up your time, reducing the amount of time you need to spend on it so you don't have to commit resources. Then, you find yourself in an implementation two years later and you don't have committed resources who use it daily or are committed to it full-time. This means you don't maintain things like the triad rules and filters. Even though the sales material says it makes it easier and reduces alert fatigue, it doesn't give more time. You still need to have a dedicated resource to operate the tool, which we never committed at the beginning.

Having an established mature team structure is really important as well. Making sure people are aware of their role and how their role fits into the use of the tool is key. Whereas, we were building a security operation center (SOC) at the same time that we took on the tool, so our analyst activities have evolved around the incorporation of the tool into the organization and it's not necessarily a mature approach.

I would rate this solution as an eight (out of 10).

Which deployment model are you using for this solution?

On-premises

© 2021 IT Central Station | www.itcentralstation.com

Read 12 reviews of Vectra AI

See All Reviews

© 2021 IT Central Station | www.itcentralstation.com