How to choose the right SIEM

Information Security controls come in two flavours. Preventative controls, such as Firewalls or authentication and authorisation systems, stop bad things happening whilst permitting the good things to go unhindered. Detective controls on the other hand tell you about the bad things when they do happen and often about the good things as well. The most notable detective control is the Security Information and Event Management system or SIEM. This device takes data feeds from all of your other security controls and probably from your IT infrastructure and applications as well. It then looks for patterns which indicate that a preventative control has failed, i.e. a bad thing has happened. Along the way it will give you pretty dashboards summarising how your controls are operating and whether you are meeting regulatory requirements (to a certain extent). Above everything though, a SIEM is there not to stop a breach but to help stop the breach becoming an incident!

The SIEM then is both an audit of your security environment and a near real time view of activity in your organisation allowing you to respond to concerning behaviours or interactions. The SIEM is arguably the most complex security control you can own given its need to interface with every other control and with the very fabric of your IT organisation. Choosing the right SIEM technology is one of the biggest decisions a CISO can make. Getting it wrong can leave you with a very expensive white elephant. Getting it right can transform your organisation and significantly improve not just your security posture but audit, risk management and even fraud exposure.

Green field or Brown

One of the most important factors in selecting the right SIEM for you is whether you are replacing an existing SIEM or deploying one for the first time. If you have a SIEM but no one uses it or if you use something that isn't a SIEM as a SIEM, I'm going to put you in the second camp here.

Replacing an existing SIEM

If you already have a SIEM then you also have a number of other things. Chief amongst these are an existing Security Operations workflow and probably existing data feeds and a data lake of some sort, consisting of some period of back data already processed by your SIEM. Both of these will have some bearing on how you proceed.

An existing SIEM workflow is a significant limiting factor in SIEM selection. That is because when you deployed your first SIEM you built up working practices and internal process models which were limited by the original SIEM's capabilities. Some of those processes probably worked around things the SIEM could not do and some were based purely upon what the SIEM did do and how it did it. Changing the SIEM forces you to consider these workflows and processes and you will need to decide how willing you are to change them. This is particularly the case if your original SIEM was deployed some time ago. You may be moving to a new generation of technology which works in a fundamentally different way. Deploying a brand new, modern SIEM and insisting on keeping the same processes is likely to severely limit the effectiveness of your Security Operations Centre - SOC. In some cases it may make your SOC non-functional if the SIEM just can't be bent to those old, embedded processes.

If you really can't change the SOC operational model, then you need to change to a SIEM as similar in functionality and approach as possible to your original SIEM. In some cases, you might consider not upgrading (assuming the SIEM is still supported and scales) but instead running a project to improve the existing SIEM deployment.

In all cases, looking at your operational SOC model at the same time as looking at new SIEM technology will significantly improve the benefits any new technology will provide. You must also however consider the adjacent technology and whether they will operate with your new SIEM or indeed if it is even needed. A good example of this might be if you include Indications of Compromise - IoC data feeds into your SIEM. Modern SIEMS almost always have these built in nowadays so having separate ones may not be necessary. You should also look at other closely coupled technologies such as SOAR. Often these tools were added on to early SIEMs and again modern SIEMs will almost certainly provide more closely integrated ones.

Clearly changing a SOAR platform will add significantly to the amount of work in your change project. But keeping an old SOAR whilst changing SIEM may be almost as much work. When checking compatibility, you absolutely must drill

into the details. A vendor saying they support your SOAR may only be part of the truth. If they don't send the data your SOAR is expecting it may not work at all or at best it may not be as effective as it once was. Problems such as compatibility will be particularly important for any incident management systems you operate. SIEM vendors will probably offer this as well but stand-alone systems still tend to be better and changing this component will absolutely force a process change on you!
You may find that you must rewrite your project and product scope after the early vendor review stages.

Another thing to think about is data retention from your existing SIEM. Depending upon your industry and what other uses your SIEM may have had over and above security monitoring, you may have a large amount of retained data. That data can be expensive to store but very expensive indeed to move. If you need to keep that data, particularly if you need to keep processing it (searching it) then you might need to run two SIEMs in parallel (your old one and your new one) for a period. You might also want to talk to your existing vendor about extracting that data in some common format, ideally into some cheap bulk storage environment. Now your old vendor isn't going to want you to do that as they will be trying to retain you as a customer, so you may need to apply some pressure even to get an answer about if this is even possible.

Older SIEMs may struggle to export the data at all, never mind in a usable format. It is worth checking if a custom solution might be built in-house to move the data. Often there are APIs available which can be leveraged for this. Another very good option is to start forking your data as you feed it into the old SIEM. Send one fork to a cheap data store and if it takes six months to replace that SIEM you will have six months of the raw data which you can either ingest into your new SIEM (probably expensive) or into some other data lake. Note that if your data is collected by the original SIEM vendor's proprietary connectors this may not be an option. You might need to install new collectors alongside those to achieve this, but if you are bound by some compliance regulation for some or all of this data you may have no choice.

One of the best aspects of moving from an existing SIEM is that (in theory at least) you should have very accurate information about your existing data volumes. If you just can't get that information out of your existing technology, then I certainly see why it's time to change! All SIEMs nowadays are priced to some extent or other on ingested data volumes. Having accurate information about your volumes is vital to build a good understanding of pricing. If you can also see growth trends for this data so much the better. But be very careful; often older (and even some newer) SIEMs do not provide accurate information. It really depends where in their ingestion pipeline they are measuring and if they are counting raw or compressed data. If the result is based upon storage data (right at the end of the pipeline) then the number may be hugely inflated as SIEMs store both the original raw events and a parsed version with particular fields indexed. The data may also be enriched by contextual data so it is not unusual for this figure to be double or more the actual ingestion figure. Or, if they are measuring the compressed stored data it might be less that half! Know what is being measured and at which point in the SIEM pipeline! Pricing will be based on data in the front door, regardless of whether you use it or store it. A good reason to filter out unneeded data from your ingestion feeds!

If you are moving from a very old SIEM to a very modern one you may also find yourself moving from a self-managed SIEM to a cloud-based SaaS SIEM. This can have huge ramifications for the entire project, not least how your data is moved across public networks, where it is stored and most importantly how it is separated from the data for other customers. You also need to understand just who can access your data if this is SIEM as a service. Ignore being pointed at SOC2 reports and the like. You need to understand who in the vendor

organisation – and/or their partners can access your data and under what circumstances. It is all very well the vendor telling you about encryption at rest and in transit, but the bottom line is – they will have to support you so they will be logging into the same (and possibly additional) screens as you and they are likely to have console access on-box for when things really break. Get them to explain and document this in detail together with all of their access and joiner/mover/leaver processes. If there is anything you can't live with or don't feel is correct then walk away. It is very likely that sensitive information will leak into these log files and having that information stored in a poorly controlled way may come back to bite you in a very big way, particularly if you reside in Europe. It is possibly time to add your data processing officer as a stakeholder here!

Green Field Deployment

Welcome to the world of SIEM. You are in for an exciting journey! The first thing to think about is what you actually need from a SIEM and this isn't as stupid a question as you may think. A SIEM is a Security Information and Event Manager. Most people tend to think of the event management piece first and foremost; information about your organisation flows in in real time, you see the bad guys and can react to them like something from a spy movie. More on this later… The other side of the SIEM is the "I" part. The information management part. This tends to be about the aggregation of data, the reporting about the operation of your SOC; about the threats you are seeing and managing in your environment; or even about the effectiveness (or otherwise) of preventative controls for the benefit of audit or even regulators.

A SIEM is there not to stop a breach but to help stop the breach becoming an incident!

You may be about to spend $Millions (or £Millions or even €Millions) on this new technology stack so at the very least you are going to want to do some reporting upwards about the use you are putting it to and the value you get from it. And a good SIEM can also significantly reduce the burden of a technology audit if you are able to produce good statistical information for your audit department at the click of a button. And in some circumstances your SIEM may also help your IT operations teams out when thing break and you can see what happened and why. Be aware though that this dual use of the SIEM may have some regulatory impacts, again, particularly in Europe.

The important thing to consider here is that you have all of the functional requirements for your SIEM thought through, with plenty of broad stakeholder engagement, before you speak to any vendors. They are going to want to talk about their new shiny buttons, which probably have AI labels on them, but you don't want to have those conversations until you have got ticks in the boxes of all of the things which you actually need.

The next thing to think about is the operational model for your SIEM. It's great having all of this data rushing into your SIEM and being analysed but if no one is looking at it them it isn't really going to give you any value for money. Don't be misled by the idea that SIEMs will do all of the work for you. They won't. A SIEM is just a tool. It is a tool which enables an operational security model in your organisation but that implies that you have people as part of that model to use the tool. Now might be the time to reach out to other similar companies who already operate a SIEM and see how they are doing things. What the operational burden is. What works for them and what doesn't. Good questions to ask are "What is your false positive rate" and "How many people do you have triaging alerts". Oh and ask them how much effort they expend upon tuning the system, particularly around detections.

And before your vendor tells you how much they have "Out of the box", particularly in terms of parsers (for different data feeds) and detections, my guess is that close to 0% of that will work in your environment without some kind of update or manual tuning and lots of it will never be relevant to you as it addresses technology you just don't have. The safe assumption is that nothing out of the box will work without considerable additional input from either your own engineers or the vendor's. And whilst we are on the subject, part of the deal must include professional services to onboard all of the feeds in your initial onboarding plan and to build out detections for all of those feeds. And ask the vendor about their training and enablement. These tools tend to be complex and you are going to need to skill-up your team. If there are any aspects of the configuration that you can't do through the GUI then for a cloud deployment the vendor is likely to charge for that work. This might be anything from adding new data sources through to data extraction and roll-off. Make sure you understand what you can do and what the vendor has to do for you.

When prioritising your requirements, remember that some things in the SIEM will need constant attention and some things will be somewhat less demanding. On boarding a new feed for instance is something you will do a lot of at the start but fairly

infrequently as time goes on. So don't focus too heavily on the tools and capabilities for doing this. If you are smart you will get a lot of free Professional Services at the start of your project and get most of this done in the first few months. Things like adding new detections will be something you do at least every month and probably more frequently. In fact every new exploit is going to see your CISO ask that "Could it happen here" question which will create a flurry of activity for your SIEM engineering team and SOC. It is probably something you want to be relatively straight forward to do, you are going to want to be able to track things like how often the detection fires and how many false positives and negatives it generates etc. This is the place you want your vendor to have spent their developers time and so think it through and have plenty of function requirements defined in terms of expected capabilities.

Which Technology?

Looking at all SIEM vendors is going to create its own set of problems for you. Originally, SIEM was a fairly simple tool, It could alert, generate a few dashboards and not much else. Search was at best limited to a few days and even that was probably theoretical. The technology didn't so much evolve as change in generations. Companies such as Splunk and Elastic arrived and the focus moved to Search as a realistic capability, often allowing you to run searches and hence SOC investigations over your entire retained data. But the fundamental architecture to do this was different from the old alerting architecture so for these companies alerting was just permanently running searches. You thus ended up with two competing approaches. One where alerting was fundamental and search was sub optimal and another which prioritised searching but where alerting could be tricky to do at scale.

These two approaches were then further disrupted by UEBA. User and Entity Behavioural Analytics. These tools started life as "SIEM Helpers" but quickly became SIEMs in their own right. They worked by using machine learning (this is just statistics done by computers, not AI, yet) to look at the behaviour of entities and spot outliers and unexpected changes in patterns. They became SIEMs when they roughly bolted on separate stand-alone data lakes to allow the search and reporting element of SIEMs. You now had three distinct approaches. UEBA tended not to do Alerting (nowadays they do) because an alert based upon changes in behaviour is even more likely to be a false positive than a "fact based" alert. Instead, they used a risk based approach which directed you to individuals users or systems and explained what you should investigate (often in the bolted on search interface).

Now the SOC workflow tends to move from detection to investigation, normally through a triage process done by first line analysts who you pay a pittance to on the understanding that it gets them into security where one day they can earn a decent living. Alert based SIEMs are great at the start of this process. Search based SIEMs are fantastic at the investigation phase. UEBA SIEMs try to reduce the triage and start the investigation phase further on than just understanding an alert. Moving from one of these SIEM types to another will require more changes to your SOC workflow than staying in your groove, and a move from a traditional SIEM to UEBA is definitely going to require a complete rethink.

One word of caution. SIEM vendors make the mistake of doing what their customers ask for. That sounds great until you remember that these vendors are supposed to be leading the way and if you ask enough people a question you will end up with every possible answer. End-users are great at saying what they want without considering the implications of getting it. UEBA vendors for instance were asked enough times for alerts that they eventually provided them. A UEBA system generating statistical alerts based upon behaviour is going to drown you in noise if your workflow is triage based.

The AI SIEM

A similar problem manifested in AI SIEMs. SIEM vendors again listened to SOC managers and heard that they were swamped with false positives and "noise"; customers desperately wanted that fixed. Instead of saying how shall we improve the fidelity of our detections they decided that they would use large language models to do the triage for SOC teams. Currently that is going about as well as you might expect. It is exactly like improving spell checkers rather than teaching children to read and write. Think long and hard about asking for AI as a functional SIEM requirement. AI is a feature, not a benefit. Ask for what you need in terms of outcomes. AI tends to ship as an add-on feature with a premium price (often subscription) and the price you pay is unlikely to go down. It is also going to have its own false positive/negative rate which will be layered on top of the SIEM's own detection capabilities. The long and the short of it is that it will become even harder to have confidence in your SIEM's ability to alert you to an ongoing breach. That doesn't mean an AI SIEM is less likely to detect a breach, just that you can be less sure that it will.

Some SIEM vendors are starting to embrace standardisation in terms of detection definitions. These definitions are open and the idea is that you can download the definition you want and ingest it into your SIEM without having to figure it all out for yourself. This is a great idea but implementation is patchy at best. To be fair, some of the problem is with the detection definition standards which, whilst open, tend to be too generic and don't really fit any SIEM technology well. They are also aimed at SIEMs which use only factual detections rather than statistical ones. They don't work so well in UEBA SIEMs.

Search

Because whatever SIEM type you select you will ultimately end-up in search. This is an area you should consider very carefully. Every investigation, every question you need to answer will be in the search interface, so you need to be comfortable with both the search capabilities themselves but also with the underlying non-functional features underpinning them.

How long your data is retained is clearly critical. If the data is nolonger there you can't search it. Many SIEM vendors have multiple tiers of data availability often conforming to Hot, Warm, Cold and Frozen status, or similar names. Hot tends to be the most expensive and Frozen often relies upon manual processes to search it at all, often taking some time to complete and on occasion not being feasible in a real-world situation. If Frozen data requires space in the Hot or Warm storage areas that may mean deleting data to allow the Frozen data to be reintroduced. It also means that you can't aggregate or report on data over very long periods. Aggregation (finding computed values like sum or mean) tends to be a very difficult operation particularly across long time periods or when spanning different data storage tiers.

Given that you will pay more for faster storage tiers, it is not unusual to see SIEM vendors only offering weeks' worth of storage for Hot and only a few months for Warm (whatever those things actually mean). Some vendors however will offer a year or more of Hot storage. The problem here is that there is no way to compare these across vendors without looking closely at the features and search performance of each tier for each vendor. Hot for one vendor may not be as performant as Warm for another.

Another aspect of the data is how enriched it is. If you have to run multiple layered searches just to understand that a username is a particular individual, what department they are in, where they work etc then your investigation is going to take a lot longer than if this information is in the original event or is injected into the results of each search. Contextual information is possibly the most valuable thing your SIEM can give you outside of the original detection capability! But beware, if the context is provided at search time it may be wrong. Perhaps that person just changed departments last week. Perhaps the IP address of their machine is different due to DHCP lease expiry. You need the context at the time of the event, not the context now for an event sometime in the past. And some of that data may time out and not be available any more.

The most obvious aspect of search is probably the search language or syntax. Some vendors have proprietary languages; some use some variation on Lucene which is an open standard. Others use SQL. Search language is often the source of online flame wars. You will invest in your SOC team learning one and despite how much they complain about it, they will resist any attempt to replace it.

The safe assumption is that nothing out of the box will work without considerable additional input from either your own engineers or the vendor's

Vendors have tried to make things simpler by allowing you to click on fields to add or exclude them. Command line completion in search bars is also common nowadays. These and other features make it easier for your analysts to get to grips with the search language with little or no training. A big new feature, again underpinned by AI, is natural language search. Just type in in English (almost always just English) what it is you want and behind the scenes your thoughts are transformed into the appropriate query. A nice feature, when it works, but almost certainly an additional cost. Do you really need it?

And what about the output of those search queries. Fine to have them on-screen when you are in the middle of an investigation but often you need to save the output as evidence, or to report your actions to your manager, or just for audit. Being able to output search results in multiple formats such as CSV or pdf is probably going to be important to you. Automated searches which publish to a URL might also be important, so that people without direct access to the search interface can still get their weekly report. Better to publish than to email – that way you can check if they are actually reading them! And being able to save

searches and share them with other users is also going to come in handy as teams work together.

A picture, it is said, is worth a thousand words. It certainly conveys complex information clearly when the picture is a chart. Dash-boarding is a feature you are going to want. It isn't going to be as useful in incident detection as you think but it is going to make your life much more bearable when it comes to reporting. Nowadays behind every chart is probably a saved search. How you build those dashboards from saved elements however is going to make a difference if you are regularly changing and updating dashboards for other teams. Don't just look at the pretty dashboard and tick the box. Get your vendor to build you a new one on the demo. If they can't do it then neither can you! Again, can dashboards be exported as static reports? Can they be made available in read-only mode to users who don't have search rights?

What your vendor will be keen to push on you are the very many "out-of-the-box" dashboards and reports. I refer you to my earlier statement about out of the box content. It is probably not worth the toilet paper it is written on. They will be keen to point out all the compliance reports. Actually stop and think about what a GDPR compliance report might look like. That's right. They have cobbled together something completely irrelevant. Most compliance regimes tend to be comply or explain. They are generally non-specific in their requirements and hence it is hard to create a report demonstrating compliance as this will vary considerably company to company depending upon what technologies each uses and how they support operational processes. SIEM Vendors want to sell content nowadays so having as many reports or dashboards or detections as possible is a key marketing requirement. Look at that content and match it to your environment to understand the potential value. And include the subject matter experts in those areas to comment upon that value assessment.

Non-functional requirements

The non-functional requirements for the SIEM are arguably just as important as the functional requirements. How the SIEM fits into your environment can make the difference between a valuable tool supporting efficient processes and something that sits on the sidelines largely ignored.

A good example of a non-functional requirement is the SIEM's authentication and authorisation model and how that fits into your broader triple A architecture. Is there a strong role based access management system for your SIEM which allows you to restrict access to data where required and share data and services where needed? Is this driven from your own Active Directory based (or other) management system or is it a stand-alone user database which is likely to get out of step as users leave and change roles? Can you generate reports for user activity within the SIEM and can you be certain that the data has not been modified, even by administrative users? This is particularly important in areas where user surveillance is not permissible, such as the whole of Europe. Proving that you aren't victimising a user may be essential in some disciplinary processes. And whilst we are on that subject, no, users do not have the right to have their activity deleted from the SIEM under GDPR so called right to be forgotten rules.

Another requirement you might want to ask about is log non-repudiation. If logs can be altered after arriving in the SIEM you have both a security problem and a regulatory one. Ask your vendor and get them to explain exactly how they achieve that. Ask them how this works when people have direct, command line access to the SIEM and the data store.

The list of non-functional requirements will be determined largely by your existing environment.

Don’t skimp on these. And again, ensure you include the relevant stakeholders. If you are deploying the SIEM on-prem or in a private cloud, there are going to be a lot more of these than if you are taking SIEM as a service.

Testing your SIEM

Once you think you have the right product and the right people to sell it to you (almost certainly not the same people) then you are going to want to make sure the thing does what it says it will, or, using the terminology of the vendor, you will run a Proof of Value (PoV) or Proof of Concept (PoC). It’s the same thing despite what they tell you so don’t worry about what they call it.

Your partner (for that is what they need to become) will want to limit the time this takes and the scope. These are good things for you also. You don’t want to tie up expensive resource playing with a new toy! Typically, they will limit you to a small number of data sources and a limited number of detections. They will want to guide you on what these should be but make sure they are ones directly relevant to you. Get the stakeholder who owns the data sources involved to ensure you can structure the feeds into your test environment.

Use a real environment if you can. It is often tempting to use a lab environment for testing new technologies but generating realistic data is a real challenge in the SIEM world. Try to use a *subset* of your production environment.

Agree the success criteria with the vendor and stick to it. They should be happy with this approach. Don’t let them move your requirements around because of their limitations. They either can do this or they can’t. Ensure the entire pipeline is included in your tests.

You may be tempted to use one of the SIEM testing companies. These systems however do not generate realistic data no matter what they say. In fact, they won’t work at all with UEBA based SIEMs which look at statistical anomalies. A single random alert from a system with no history is not a realistic measure of a UEBA system’s capabilities. The bottom line is that whether you use a RED team or an automated evaluation system, the partner will try to tune for that. That is they will create detections looking for those specific events. The problem is that that gives you no insight into how the SIEM will function generally and how much noise it might generate in production. Give them no information about how you propose to test this PoV. You need to understand the realities. And most importantly, pay close attention to how they configure the PoV. Make sure your security engineering team are on every call and watching every screen. Is this difficult for them or just a few clicks? Is this really out of the box or does everything need to be reworked?

Specialist SIEM features

One great feature, which is becoming more prevalent is data routing. This can be particularly valuable when your SIEM separates the core functionality from the search capability and charges for the former but not the latter; or where the analytics part, normally of a UEBA SIEM is more expensive than the data lake element. The ability to route data to the right place and often to filter out data entirely, that doesn’t need to go to the SIEM can save you a considerable amount of money. Also, you might value the ability to duplicate flows and send both to the SIEM and to other systems. This technology has been available for some time both as commercial stand-alone products and as open-source capabilities but increasingly it is built into SIEM systems. Decide first if you need it then if it is better to have it within the SIEM or separately. The chief question is what collects your data from the originating systems?

Another useful features is *variable roll-off* of data. Keeping some data sources in the data lake for long term search whilst timing out the rest to save money or maintain performance. For vendors that charge you for total retained data this can be really important, particularly for industries where retaining some data is a regulatory requirement. A lot of the data you collect will be of limited value whilst some other data sources may give tremendous value and be useful to search for many months or years. Generally, the further up the stack you go the more rich and valuable the data is. Network traces are often useless except when fault finding a live situation where as application logs may hold context that stays relevant for years.

More recently, SIEM vendors have started to offer customers the option of using their own data lake at the back end. The search interface remains the SIEM, though users normally have the option of also leveraging their own native tooling to interact with the data. Be warned, this is rarely a cost saving. The chief benefit of this approach is enhanced control of your data and this can be very important where you want very long retention periods or where you need to retain all data after changing vendor. Don’t even think about signing up for this until you have spoken to your data team

about it. They will want to understand both data volumes and access patterns.

SIEM add-ons

SIEMs are no longer stand alone. They generate alerts which often get promoted to be incidents and these tend to be managed in incident management systems. Most SIEMs come with Incident Management but these can tend to be both opinionated and limited. An incident often requires input from non-security teams so a SIEM-only solution might be too limiting to allow this to work effectively. Most SOC teams end up using third party incident management systems which are generally more capable. Using third party systems also means being independent of the SIEM if you choose to change that in the future. If you plan to use, or already have, a third party Incident Management system, ensure your SIEM will provide effective two way interfaces to it. You need data to pass from the SIEM into your Incident Management System (IMS) and back again, to cancel incidents etc.

SOAR

One particular SIEM add on is the SOAR, or Security Orchestration, Automation, and Response system. When first invented, security vendors envisioned a world where every detection would reach back into your IT environment and adjust configurations, isolate systems, change Firewall rules etc. Clearly these people had never worked in a real organisation with separation of duties, change management systems and IT governance. Nowadays SOAR does somewhat less than was envisioned and is much more of a pain to maintain than was hoped. Typically there are a very few really good use-cases for SOAR. Check the Vendors one meet those requirements and can be expanded upon, ideally with python or some other common scripting language.

Don't ask vendors for features no SIEM has. It is a waste of everyone's time, including yours. It doesn't matter how much you want a pink, fluffy, unicorn, if there are none to be had, that is that!

Every SIEM has a SOAR add on option. Some were built by the SIEM vendor and some were acquired.. There are only a few stand alone SOAR vendors so your best bet is almost certainly to take the SOAR that comes with the SIEM. Interfacing external SOAR systems can be a real pain. The rule here is to think about exactly what you need your SOAR to do. Many things which were previously done with the SOAR are now built-in. Some companies use SOAR to export incidents to IMS systems. Again, you may or may not need that. Make no mistake; a SOAR can be a lot of work for your security engineering team. It generally requires a good understanding of both the SIEM and the target system the SOAR needs to interact with to build a new interface. Your team may need some programming skills to pull that off effectively. And every time either of those interfaces change you'll be doing it all again.

The best advice is: get your SIEM up and running and then come back and look at what SOAR might give you. That tends to be less than you think. Try to ignore the nice user interface where the vendor demonstrates drawing a flowchart and the SOAR then magically does all of those things, talking to three or four external systems. If that happens like that even once in your environment then you are winning. Hint: It won't.

IOC Feeds

Indications of compromise or other data feeds into the SIEM are a popular add-on. The idea is that this enriches the detection capabilities of the SIEM. Most vendors will have a free feed or two which is probably one of the open source feeds supplemented by their own "security researchers". You can also choose to buy variously priced commercial feeds and even industry specific feeds. The reality is that these data feeds provide fairly limited value. Often SIEMs have better ways to do the same job such as identifying new or malicious domain names. The reality however is that everything outside your organisation is either bad or probably bad. Knowing that one PC is sending large amounts of data to a known bad IP address is only very slightly better than knowing that one of your PCs is sending large amounts of data out of your organisation. Are you really only going to investigate when it is a known bad IP address or domain? As always it comes down to your workflows and what your triage process looks like.

Some organisations have more than one SOC. It may be that way due to years of mergers and acquisitions or because you are a global organisation or split by business unit. Either way you may wish to have each SOC able to access its own data whilst a corporate level security team has visibility over the entire business. Choose a SIEM which is designed for Managed Security Service Providers (MSSPs). These tend to have a tiered system but with the ability to share data connectors

and detections across sub organisations, as well as very advanced AAA controls (Authentication Authorisation and Accounting).

Top Tips

Ensure you involve all stakeholders. Too often have I seen security engineering teams selecting a SIEM with little or no engagement with the SOC manager.

Where is your data? If it is already mostly in a cloud provider's lair then having a SIEM in that same place probably makes sense. Sending data from one cloud provider to another can be expensive. Check out cloud data exfiltration costs. If your data is mainly on-prem but you have a plan to move to the cloud then a cloud based SIEM might still be right for you. If you are a global company, particularly in certain industry verticals, you might find that some data can't move region. In some cases it may not even be acceptable to access that data from another region. You may be forced to have multiple SIEMs in that case. This tends to reduce visibility into global threats but it may be unavoidable.

Your SIEM project may be the largest security project your company has ever undertaken. You won't get it done in a week. Or even a month. In fact in year two and three you may still be bringing in data feeds and you will certainly still be configuring new detections. Your security engineering team will be as busy as your SOC team so plan for that. Also ensure that you have great support and ideally recurring Professional Services time available to you for the duration of your contract.

Changing SIEM hurts. Make sure that your vendor guarantees that the SIEM will be supported through the entire life of your contract and beyond.

Don't buy road-map! It is either a feature in the main product or it does not exist. Signing up with a vendor because of what they say they are working on is a fool's errand. Don't do it. And be oh so careful. Some vendors will clumsily build features in the field for your demo environment which are not in the core, supported product. Record the demos and get confirmation in writing that every feature discussed or seen are in the core product in the form you saw it.

Everything you didn't see in the demo but are promised needs to be dealt with also. If the vendor assures you they have support for your Incident Management System that doesn't mean that they support it the way your current SIEM does or the way you expect it to be done. Get it specified in as much detail as possible and get it signed-off.

After signing on the dotted line you will never see the vendor team you have been working with again. Instead you will suddenly find yourself talking to new people responsible for the post sales deployment and support of your new technology. Do yourself a favour and ask to meet those people well ahead of any agreement being signed. Check that they have been well briefed and talk to them about how they engage with customers. Better yet, tell them you want to build a mutual success plan with them before signing.

On the subject of signing, find out when your vendor's end of year is. If you can take signing up to the wire then they will get increasingly desperate and are more likely to offer you more concessions to get you to sign. Much better to get them to offer you more Professional Services or a higher tier support contract than to actually drop the price. They are normally paid on the actual licence quota rather than on these extras so they will happily give you lots more of that to close the deal.

Beware of hidden limitations or costs. Ask if every feature is unlimited, such as the number of alerts you can configure. This is particularly the case where the SIEM is hosted in the vendors cloud. They will seak to protect their own cost by forcing you to adhere to some notional acceptable use policy.

Lastly, don't ask vendors for features no SIEM has. It is a waste of everyone's time, including yours. It doesn't matter how much you want a pink, fluffy, unicorn, if there are none to be had, that is that! And that goes for features that don't belong in a SIEM, like Vulnerability Management. It might exist but if you need that feature then go and buy that product.

If the one feature you really need is only provided by one vendor then asking all of the other vendors for it will just tell them that you have already decided who you are going with and that will prevent them from taking you seriously. Similarly, if you really want a UEBA SIEM then just send your Request for Proposal (RFP) to UEBA vendors, not to every little SIEM player out there. Vendors will look at any new RFP and decide if it is open, written for them or excludes them. Unless they are really struggling they will no-bid (ignore or decline to participate) your RFP process.

You are going to have to live with this product and this vendor for a few years. Choose wisely!