securonix

Cyentia INSTITUTE

FINDING THE SIGNAL THROUGH THE NOISE

Quantifying SIEM Effectiveness

Contents

Executive Summary

For this report, Securonix partnered with the Cyentia Institute to analyze a dataset of more than 54 billion events fed into more than 154k policies generating an average of more than 750k violations per hour. The goal? To quantify our assumptions and findings in a way that can help organizations calibrate what's going on in their own environments.

More data sources will beget exponentially more policies. We found that two-thirds of organizations utilize between seven and 17 different types of technologies to monitor their networks. For approximately every seven data sources an organization adds to their repertoire, the number of policies doubles, which means that many organizations quickly find themselves with hundreds (and sometimes thousands!) of policies worrying away at all the events within their environments.

More visibility means more signals to analyze. More policies generate more signals. In fact, when an organization doubles their policies we see an approximately 6x increase in the amount of first layer alerts. That means six times as much that needs to be analyzed, either manually or by a second layer of analytics.

More signals to analyze correlates with less signal actually analyzed. We see a 42.2% decrease in the number of adjudicated violations when an organization's violation rate doubles. This could mean that the additional information is providing the context needed to focus on what's important or that things are becoming overwhelming, either way prioritization is key.

Prioritize signal, rather than focusing on increasing noise. Creating and maintaining "good" policies will allow your organization to see the most impact on adjudicating concerning violations, while dismissing the non-concerning alerts. What data source an event comes from and how it's monitored affect the quality of the signal. Having the broad knowledge base of a cloud-based Next-Generation Security information and event management (Next-Gen SIEM) can reveal more than trying to work in isolation.

Introduction

The phrase “what you don’t know can’t hurt you” has never been more wrong, than when it’s used in security. More often than not, it is exactly the things that we **don’t know** about that attackers can end up using to sneak into networks. Even if we interpret the phrase to mean “ignorance is bliss”, this aphorism still fails as it’s all those ‘unknowns’ (both known and unknown) that keep most security professionals up at night.

Of course, we are far from the first to make this observation. As an industry, we’ve worked hard to gain better visibility into our environments. First, we began logging to create a history of what processes and network connections were doing throughout the day. Then came the SIEM, which actively monitored those logs as they were being created and, of course, if we could monitor networks, we could also monitor all kinds of software and devices. So once those logs became voluminous enough, clever rules were created to narrow the scope to focus on the important stuff. But eventually, even that filtered data became un-monitorable by bedraggled SOC analysts, and more advanced analytics were born to help not just separate the wheat from the chaff, but also connect and aggregate findings. Finally, if we could automate the collection and analysis of events, we could also automate the response, right? And just like that, we moved into the realm of Next-Gen SIEM.

Whereas not knowing was the liability before, now most organizations are cursed with more knowledge that they can reasonably process into meaningful action. This overwhelming vision of an organization’s environment, and all its goings-on, has led to the obvious conclusion that “there must be a better way”; there must be patterns in all this data that we can have machines recognize and focus our vision on what really matters.

In the spirit of gaining a better understanding of just how organizations are handling this flood of events, Securonix partnered with the Cyentia Institute to analyze a sample dataset of more than 54 billion events fed into more than 154k policies generating an average of more than 750k violations per hour. In particular, we wanted to see what organizations use to monitor their networks, and take a look at what provides useful, actionable information.

Some of what we found didn’t surprise us; the more things you choose to monitor, the more noise you’ll have to sort through. Some systems produce real, important information on a regular basis, while others...not so much. But this report will do more than just point out the obvious; it will quantify the obvious in a way that can help organizations calibrate what’s going on in their own networks. Too often in security, we are restricted to our own (or our organization’s) view of the world, which makes it impossible to find larger patterns. Securonix has a unique view here; we can aggregate information across a multitude of customers, allowing a large enough sample to be able to better understand what works and what doesn’t.

Some terminology

There is no common set of terms surrounding next generation SIEM technologies and different vendors typically use different terminology, so we start with a few brief definitions to help aid the reader in understanding our findings.

Alert

A violation that is brought to the attention of a security analyst and adjudicated by them. Not all violations rise to the level of needing human intervention and many policies will generate violations, but never alerts.

Concerning/ Non-concerning Alert

Indicative of whether a human has decided that an alert warrants further investigation (concerning) or not (non concerning).

Data Source

This is a piece of software or hardware that generates events.

Event

A small piece of information that gets conveyed to a Next-Gen SIEM.

Policy

What are generally referred to as "rules", but encompasses more than simple SIEM rules. Policies can range from a simple bit of logic to complex machine learning analytics executed by the SIEM on incoming events. Policies describe to the SIEM what conditions in the ingested events should generate violations.

Policy Type

A categorization of various policies.

Violation

A piece of information generated by a policy after seeing events that meet a particular condition. Violations can indicate potential malicious behavior or can simply provide context for security analysts.

What are y’all watching?

Perhaps the first question we might ask is “What exactly are organizations monitoring?” With the multitude of technologies out there meant to make our everyday work easier, there are a lot of places we can look for malicious behavior. Figure 1 shows the popularity of monitoring different data sources.

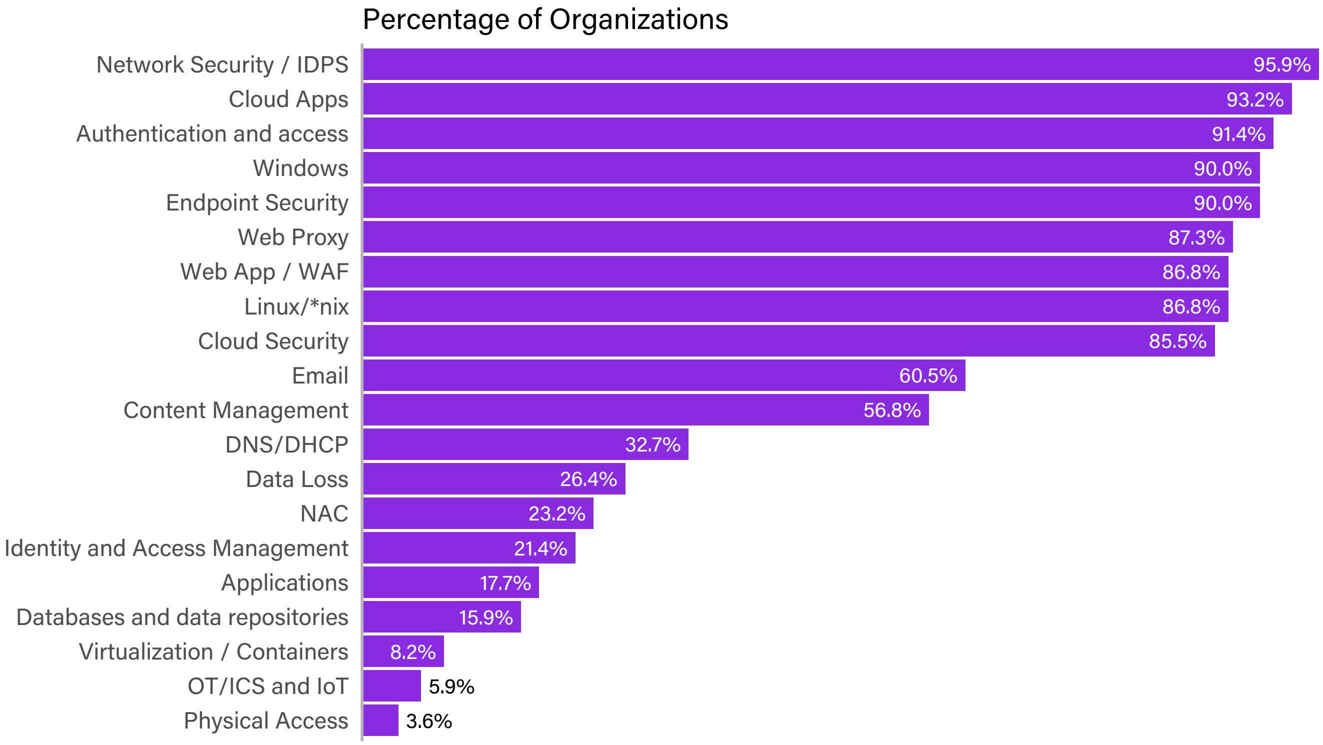

FIGURE 1 – DATA SOURCE POPULARITY

There are of course caveats to results like the above. Not all products do exactly one thing and terms tend to be a bit squishy. One vendor’s authentication monitor might handle Cloud, SSO, and local authentication, but others might only monitor local authentications. Another caveat is that not all organizations are likely to dump all their data into SIEM[1]. It’s possible (or likely) that nearly every organization monitors some of the data sources lower in Figure 1 (I’m lookin at you IAM), they just aren’t feeding that particular data source to Securonix. But we think few would argue with the relative order and prevalence here; for example, given that SIEM was originally designed for network monitoring it’s unsurprising it tops the list.

[1] Securonix

Now we know roughly from what types of things are being monitored, we can also ask: “How many different types of things are being monitored?” Figure 2 has the answer:

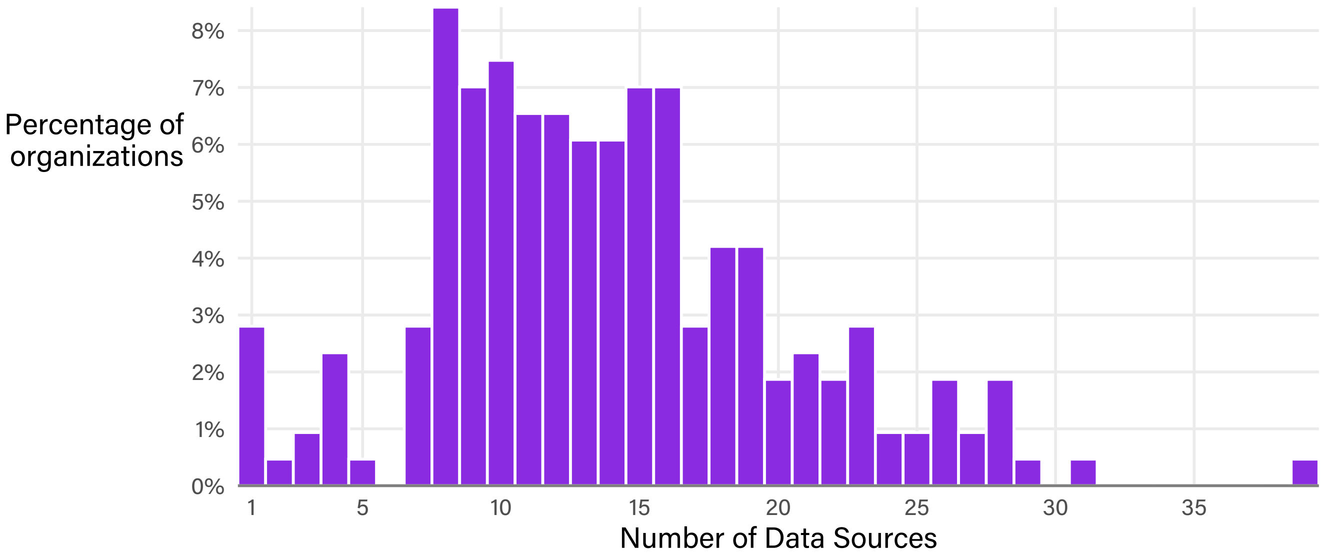

FIGURE 2 – NUMBER OF DIFFERENT DATA SOURCES

2/3rds OF ORGANIZATIONS UTILIZE BETWEEN 7&17 DIFFERENT TYPES OF TECHNOLOGIES TO MONITOR THEIR NETWORKS

ONLY 7% OF ORGANIZATIONS MONITOR <7 SOURCES

THE TOP ORGANIZATIONS WATCH >30 DIFFERENT TYPES OF DATA POUR ONTO THEIR PLATFORM

What is interesting here (and will become more obviously so in later sections) is the relatively small range of data sources. We’ll see later that by any other measure organizations differ by 10x, 100x, or even 1,000x. All that variation comes from a relatively small set of data sources.

So how does this small set of data sources expand? Read on...

It's not just what organizations monitor, but rather *how* they monitor it. That is, it's not enough to have an appliance or software feeding data into a platform, you also have to write policies that check for patterns in the data that indicate when something isn't quite right. So, we wanted to find out how these policies proliferate as more monitoring functionality is added to the platform.

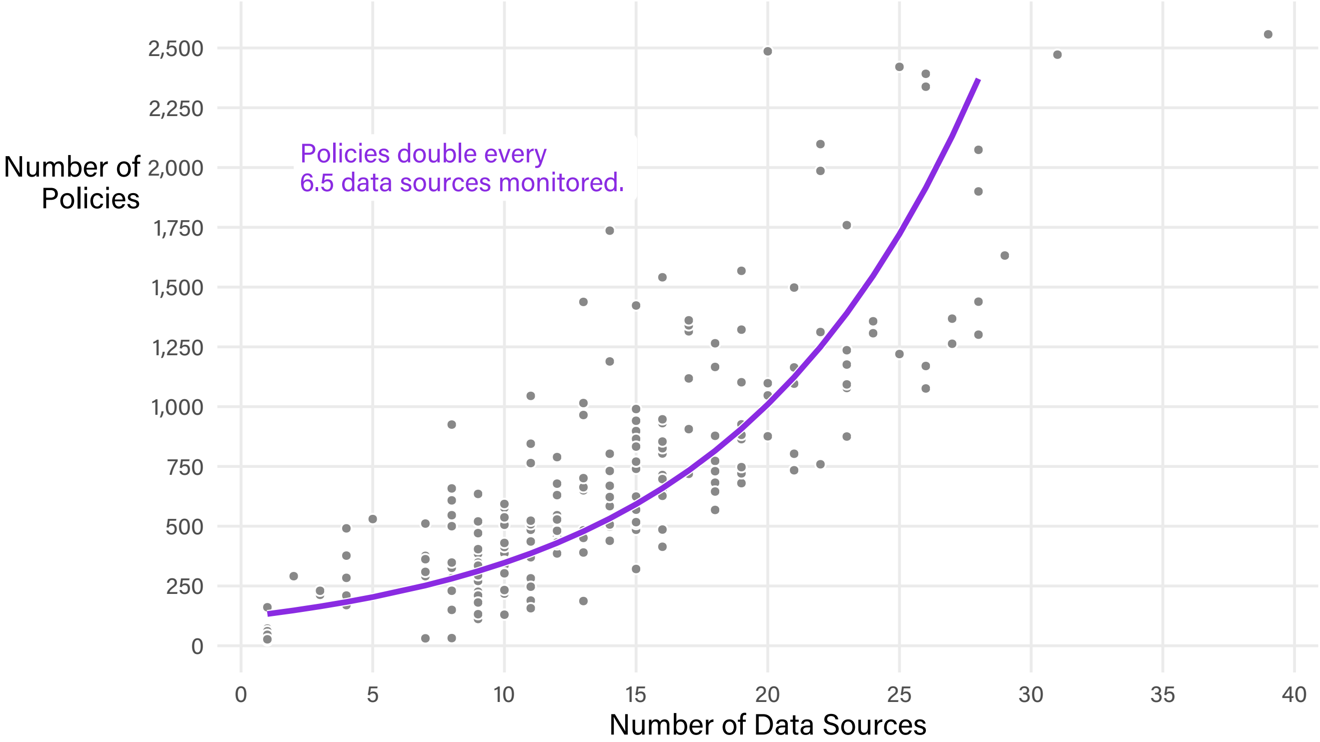

FIGURE 3 – ORGANIZATION POLICIES VS DATA SOURCES

What is surprising about Figure 3 is not necessarily the fact that as an organization adds more data sources that it deploys more policies, but rather that the number of those policies grows *exponentially*. In particular, for every seven data sources an organization adds to their repertoire, the number of policies doubles. As can be seen in Figure 3, that means that many organizations quickly find themselves with hundreds (and sometimes thousands!) of policies worrying away at all the events on a network.

As we ponder the implications of Figure 3 and look ahead to Figure 4 and 5, we should mention that these findings may be directly related to the characteristics of the Securonix product. The number of policies added per data source, for example, can be different for products that provide less "out of the box content". As many organizations use the content provided by Securonix, the distribution of techniques may also be related to characteristics of our (Securonix's) content, instead of what organizations are writing themselves.

It's important to consider not just how many policies organizations have in place or what those policies are monitoring, but also exactly what techniques a policy uses to turn the data it sees into actionable information. This can range from simple, quickly evaluated rule-based policies to advanced analytics using the latest machine learning techniques. In Figure 4 below, we look at what types of policies organizations deploy.

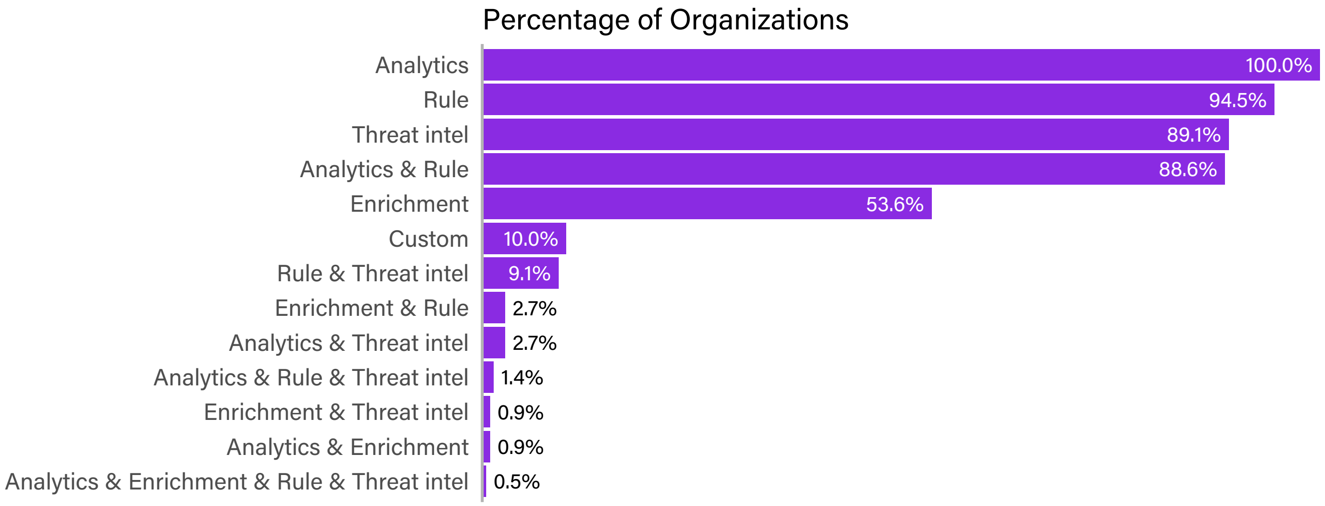

FIGURE 4 – POLICY TYPE POPULARITY

The vast majority of organizations use policies that utilize simple rules, advanced analytics, and threat intelligence. About 50% of organizations use "enrichment based" policies, where data is pulled from other sources to check, for example "is this user logging in from a geographic location a long way from where they recently logged ?". It's worth mentioning that Enrichment policies depend on enrichment data being available, and many organizations are still not mature enough to add a substantial number of those. We should see more enrichment policies as organizations integrate more context sources to their SIEMs.

Policies can of course leverage more than one technique. A large number are also combining techniques and constructing policies that trigger violations when a combination of data sources are used, though rarely more than two. Perhaps more interesting is on what types of data sources different types of techniques are used. Figure 5 breaks it down.

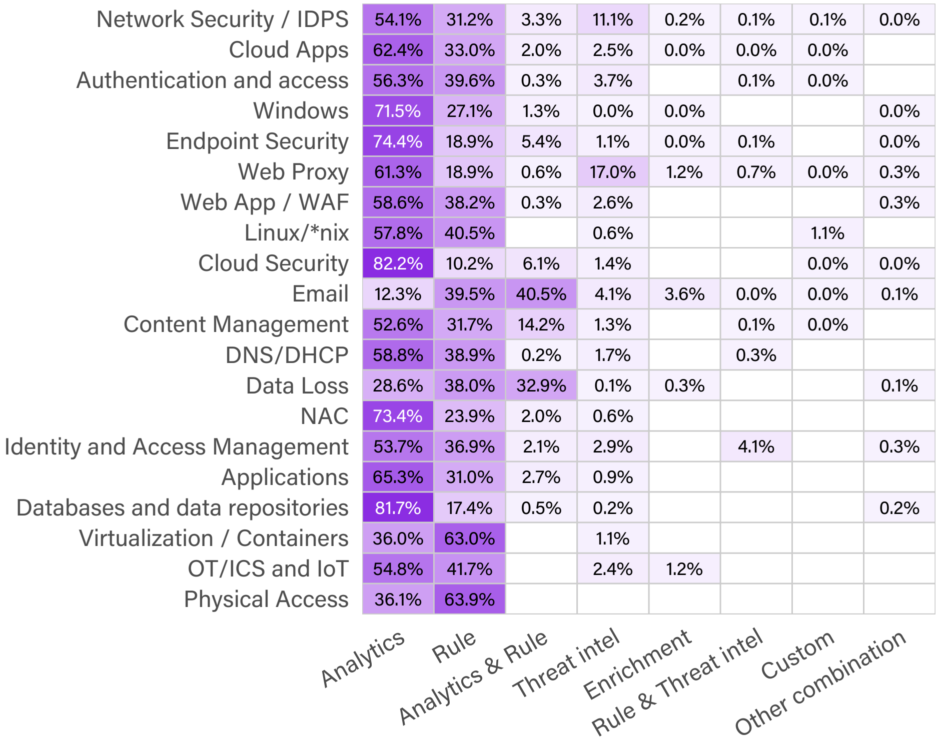

	Analytics	Rule	Analytics & Rule	Threat intel	Enrichment	Rule & Threat intel	Custom	Other combination
Network Security / IDPS	54.1%	31.2%	3.3%	11.1%	0.2%	0.1%	0.1%	0.0%
Cloud Apps	62.4%	33.0%	2.0%	2.5%	0.0%	0.0%	0.0%	
Authentication and access	56.3%	39.6%	0.3%	3.7%		0.1%	0.0%	
Windows	71.5%	27.1%	1.3%	0.0%	0.0%			0.0%
Endpoint Security	74.4%	18.9%	5.4%	1.1%	0.0%	0.1%		0.0%
Web Proxy	61.3%	18.9%	0.6%	17.0%	1.2%	0.7%	0.0%	0.3%
Web App / WAF	58.6%	38.2%	0.3%	2.6%				0.3%
Linux/*nix	57.8%	40.5%		0.6%			1.1%	
Cloud Security	82.2%	10.2%	6.1%	1.4%			0.0%	0.0%
Email	12.3%	39.5%	40.5%	4.1%	3.6%	0.0%	0.0%	0.1%
Content Management	52.6%	31.7%	14.2%	1.3%		0.1%	0.0%	
DNS/DHCP	58.8%	38.9%	0.2%	1.7%		0.3%		
Data Loss	28.6%	38.0%	32.9%	0.1%	0.3%			0.1%
NAC	73.4%	23.9%	2.0%	0.6%				
Identity and Access Management	53.7%	36.9%	2.1%	2.9%		4.1%		0.3%
Applications	65.3%	31.0%	2.7%	0.9%				
Databases and data repositories	81.7%	17.4%	0.5%	0.2%				0.2%
Virtualization / Containers	36.0%	63.0%		1.1%				
OT/ICS and IoT	54.8%	41.7%		2.4%	1.2%			
Physical Access	36.1%	63.9%						

FIGURE 5 – DATA SOURCES MONITORED BY DIFFERENT POLICY TYPES

In Figure 5, to the left, each row represents the percent of policies that monitor a particular data source with a particular technique. As we can see, analytics dominates most categories, but there are some exceptions. Containers and Physical Access are still largely monitored by rule based policies by a factor of 2 to 1. This is primarily the result of users bringing relatively simple policies such as "generate a violation when this container crashes" and "whoopsie, this badge reader became disabled". Sometimes the simplest rules are the best. For email, just analytics aren't good enough, with 80% of policies using Rules, either in isolation or in combination with Analytics.

Having a lot of policies leveraging all the latest techniques, ready to alert you of any potential malfeasance is fantastic from a threat coverage point of view. If any of those policies are triggered by real badness, we'll be glad they're in place. But the flip-side of that threat coverage coin, as nearly all readers of this report are aware, is that many of these policies don't generate information that is at all useful to humans. When we look at the volume of alerts being received, it can sometimes feel like standing on the receiving end of a firehose. So, how can we, and organizations in general, start to understand what streams of information are important for what we care about? In the next section, we'll look at how fast those alerts are generated, and how many people pay attention.

Violations: Fast and Furious

Customers monitor a handful to a few dozen data sources with an exponentially increasing number of policies. But what does the stream of events from those data sources and the violations created by the policies look like? Let's take a look at Figure 6.

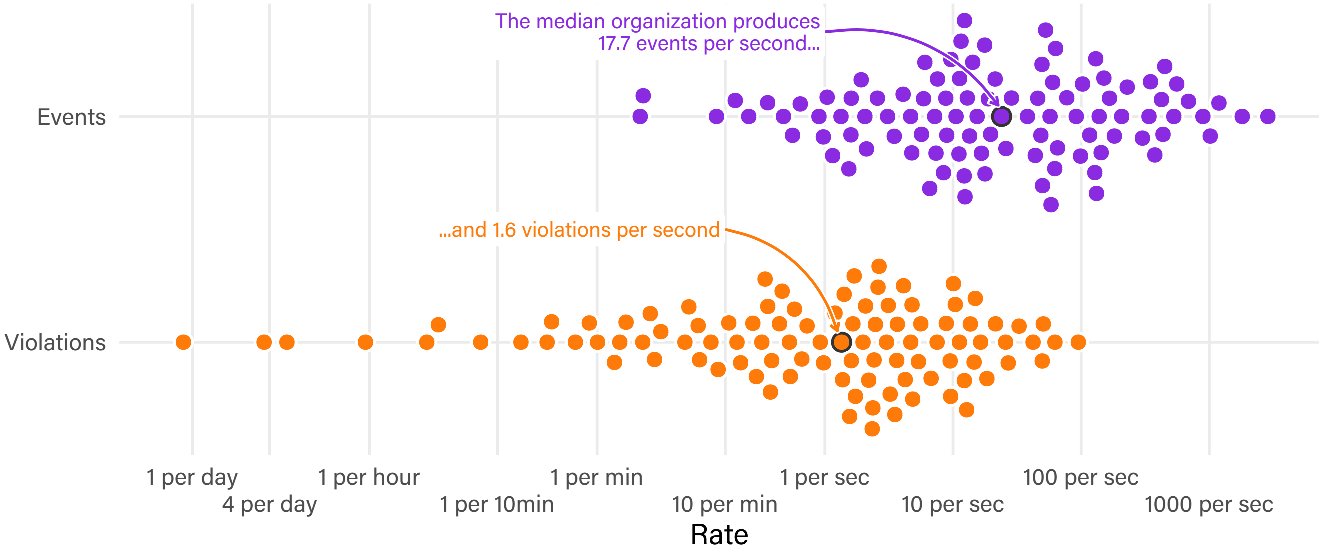

FIGURE 6 – RATES OF EVENTS AND VIOLATIONS

Each dot in Figure 6 represents approximately 1% of organizations. Each dot is placed horizontally at the measured rate for that 1% of organizations. Dots are stacked to avoid each other and create wider stacks when more organizations have a similar rate.

We can see that organizations are ingesting at *least* one event per minute, with a few experiencing more than 3,000 per second[2] (note the log horizontal scale). This funnels down to less than 100 violations per second, with most organizations in our data set seeing more than one per second. Looking at the figure, we are seeing expected behavior - especially when we funnel from events to violation.

[2] Note this analysis was conducted on only a sample of representative customers. The long tail of "events per second" extends well beyond 3,000 per second.

As a quick aside, it's worth clarifying that violations are usually not what is presented to the analyst, and the "events to violations" ratio is just the first layer of the events to human eyes funnel. There are other reductions based on combinations of violations and additional analytics to further reduce the volume of information provided to the analysts.

OK, BACK ON TRACK, SO, WHAT CAUSES A VIOLATION, AND WHERE DO THEY ALL COME FROM? THE POLICIES, OF COURSE!

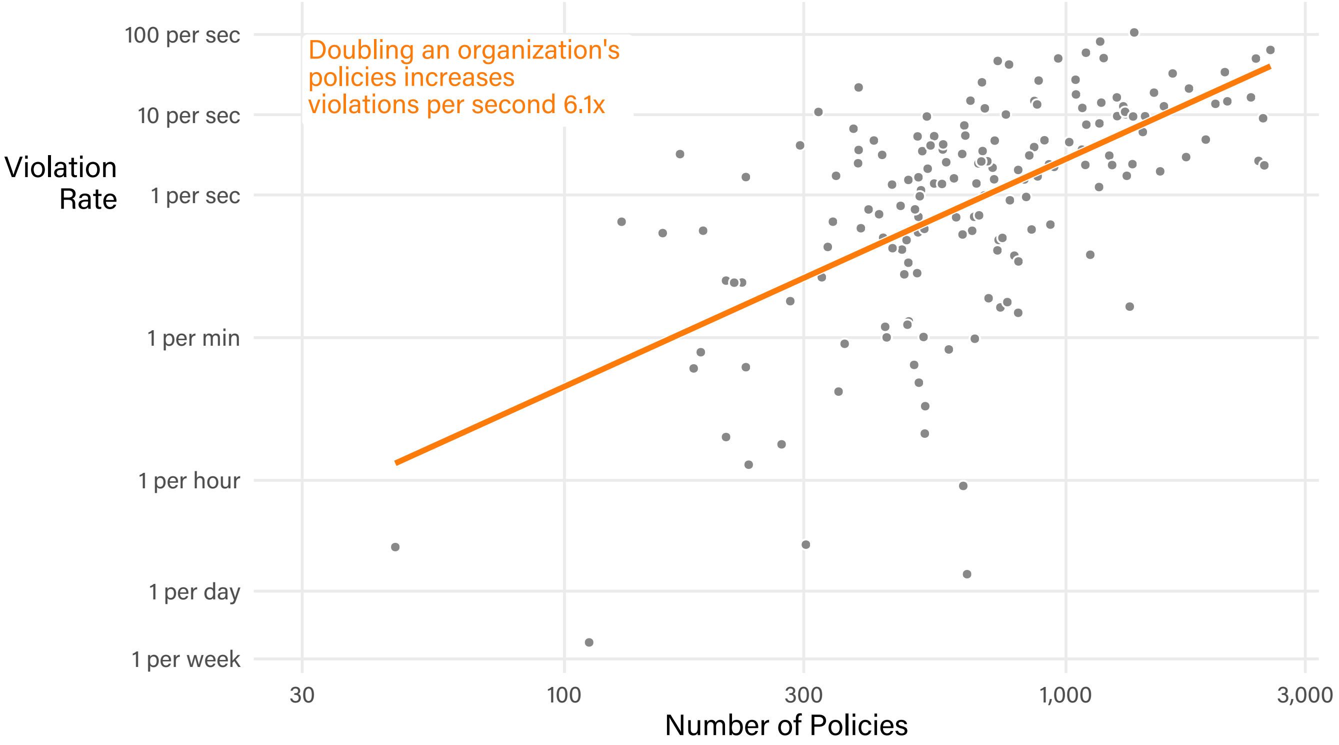

FIGURE 7 – VIOLATIONS GROW WITH POLICIES

2x POLICIES =6.1x INCREASE IN VIOLATIONS PER SECOND

Taking a look at the Figure 7, we can see pretty clearly that the violation rate increases faster than we might expect it to. We know that when more data sources are monitored there are more policies. When there are more policies, then there is the added effect of having a higher violation rate, because there are more instances when a "violation" is possible. When an organization doubles the amount of policies they have, there is a 6.1x increase in the number of violations per second.

As we mentioned in the brief aside above, violations are not necessarily meant to be delivered to the analyst. They are meant to be building blocks that create a contextual picture of potential malfeasance that an analyst can, well, analyze. This is of course how things should be. One violation per second would be overwhelming for manual analysis, but is trivial for the right set of algorithms. So the next big question: How many violations are raised to the status of "alert" and adjudicated by an actual human? Figure 8 says "not a lot".

1 VIOLATION PER SECOND

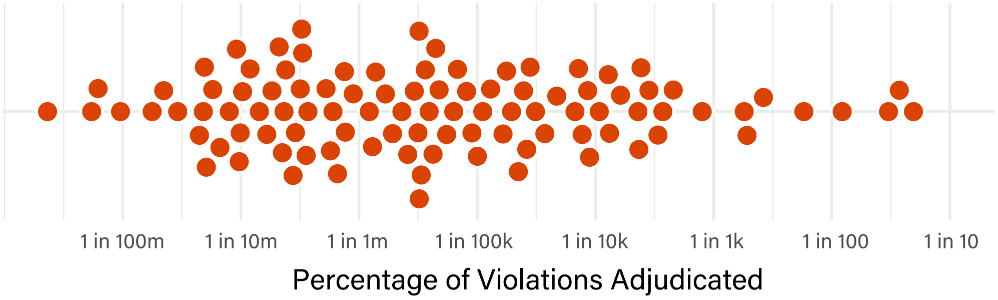

FIGURE 8 – DISTRIBUTIONS OF VIOLATIONS ACTUALLY ADJUDICATED BY AN ANALYST

The big takeaway here? For every organization in our data set, the number of alerts to analysts is 10x lower than the number of violations, with most organizations examining fewer than 1 in 100,000. This might feel like a bad thing, but as we said it's not. We promise.

THERE ARE FOUR MAIN REASONS WHY:

1 Many violations are transitory. Who hasn't attempted to login from a strange network while traveling? Or accessed a website that happened to be included in a threat intelligence feed? All those policies will raise a violation, but they don't always mean something malicious is happening. A conservative strategy makes note of everything that could be bad, so when things get really bad, there is plenty of context to draw from. And modern SIEMs (like Securonix!) look at multiple signals in aggregate before deciding to alert an analyst.

2 Some violations are used to increase the risk score of entities, or flag something for your watch list. They are meant to build context and provide a degree of severity rather than be alerts in and of themselves.

3 Different organizations will have different philosophies and different strategies for these events. It will most likely be tied to their resources and "normal" violation pace.

4 Finally, organization practices are likely different, some organizations might examine violations but never note them as having been examined. It's definitely possible that organizations are looking at 10x or 100x more violations, but aren't bothering to mark them as such. Others might be fastidious and ensuring everything an analyst looks at gets noted.

But does the firehose of violations affect the number of generated alerts?

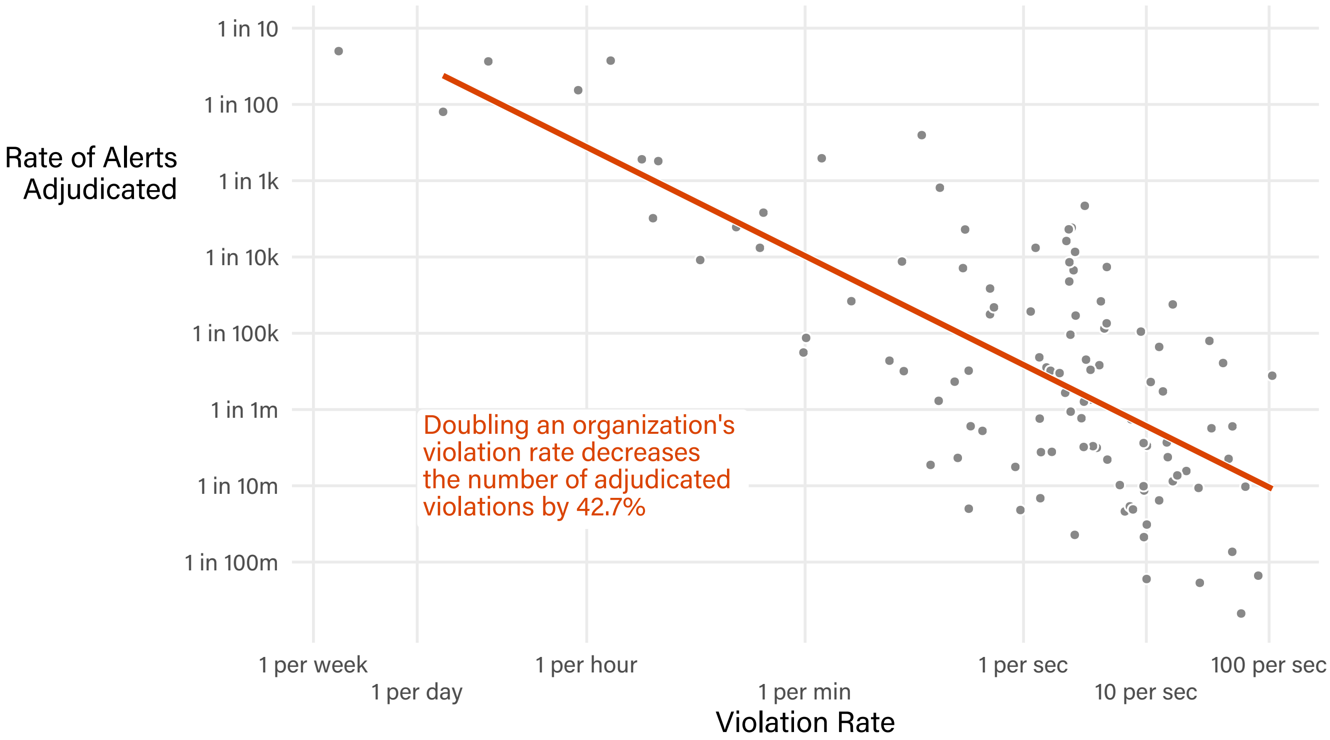

FIGURE 9 – RELATIONSHIP BETWEEN AN ORGANIZATION'S ADJUDICATION RATE AND VIOLATION RATE

Figure 9 begins to get to the bottom of it. When we look at the violation rate against the adjudicated rate, we can see a lot of variation going on. As we noted previously, a lower adjudication rate is not necessarily worse, since we don't know what policies a particular organization has in place that may trigger a violation. What we can say here is that doubling an organization's violation rate *decreases* the number of adjudicated violations by 42%.

Two possible reasons for this are the following: 1) More violations are providing additional context so fewer, more precise alerts are generated or 2) organizations are raising the threshold of what is considered worrisome to keep up with the volume of data. The former is good, the latter... less so. For each organization it's probably some combination of the two and a balancing act between their threshold for risk and the resources available.

One last bit of information for this section that might be useful to some readers. Exactly how quickly are organizations investigating alerts.[3] We see this distribution in Figure 10.

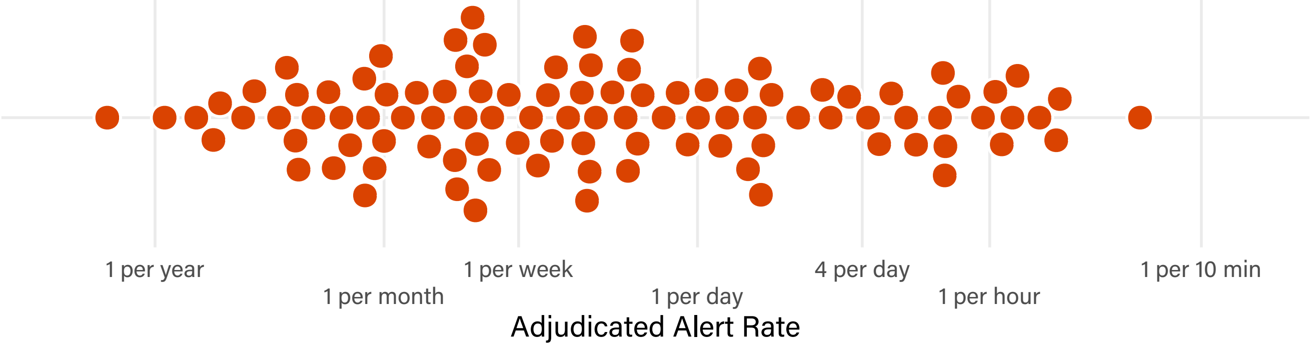

FIGURE 10 – DISTRIBUTION OF ALERTS ADJUDICATION RATE

[3] This is the product of the two values in Figure 9.

One one end, there are organizations that are investigating and adjudicating a few dozen violations a day, while others are on the complete opposite end of the spectrum, touching just a handful a year. The above figure on its own could be the start of its own research report: “Why do some organizations examine so many more alerts than others?” The answer probably lies at the nexus of organization size, mission, and security culture, but lies beyond the scope of our current research.

Seeing if there is any “there” there

The next stage in an alerts journey to the attention of an analyst is the decision of whether an alert, when finally adjudicated, is actually an indication that there is a “there” there. This journey is a long one, and only a minority of policies generate data that are ever marked as “concerning”. Figure 11, shows this funnel of policies.

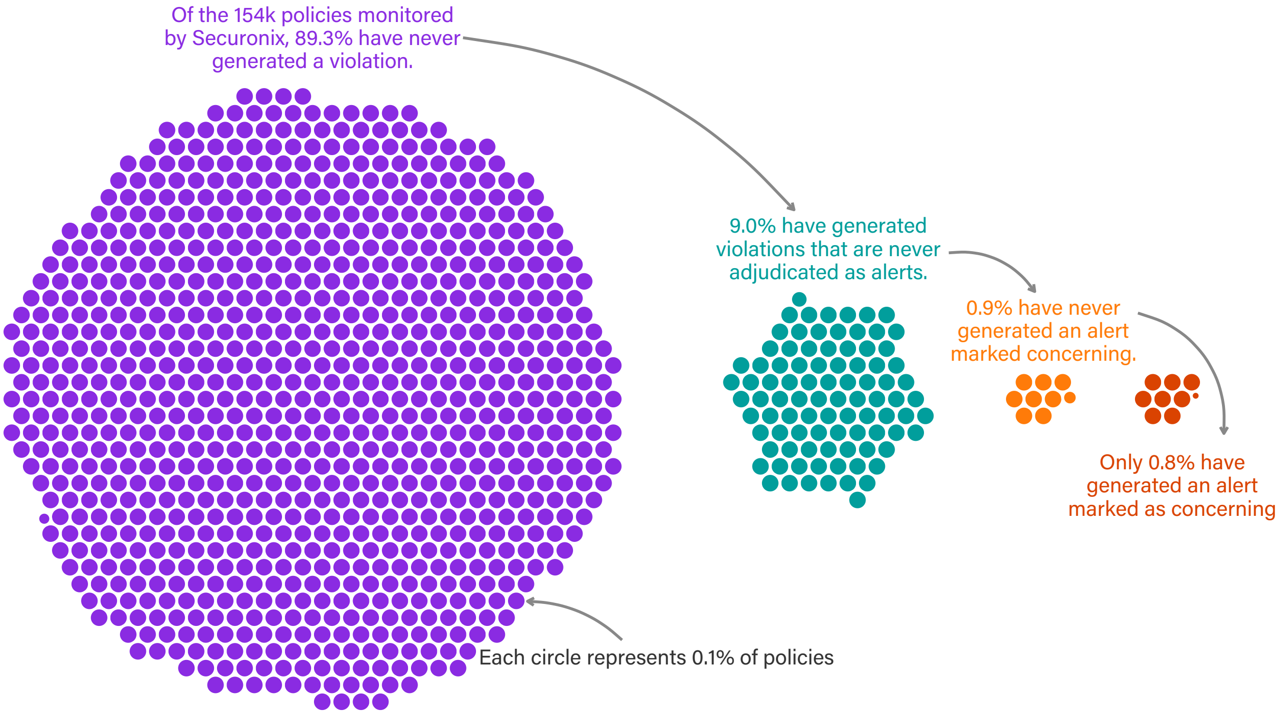

FIGURE 11 – THE FUNNEL OF POLICIES

We’ve talked about the first three stages of this funnel, but not the last one; it helps to see them all in one place. A little less than 1% of policies are doing the work to generate something, anything that an analyst would find concerning. These are what we want to focus on in order to see what is providing the really strong signal of malicious behavior. Again, this does not mean the other 99% of policies in place are bad—they are not, and the context they provide can be extremely useful.

Up until now, we’ve mostly focused on findings from the organizational level. However, it’s worth it to start to still down into individual policies to start to see a clearer picture on what is actually going on. Specifically, we know some policies are much chattier than others, so let’s quantify what “chattier” actually means when it comes to alerts.

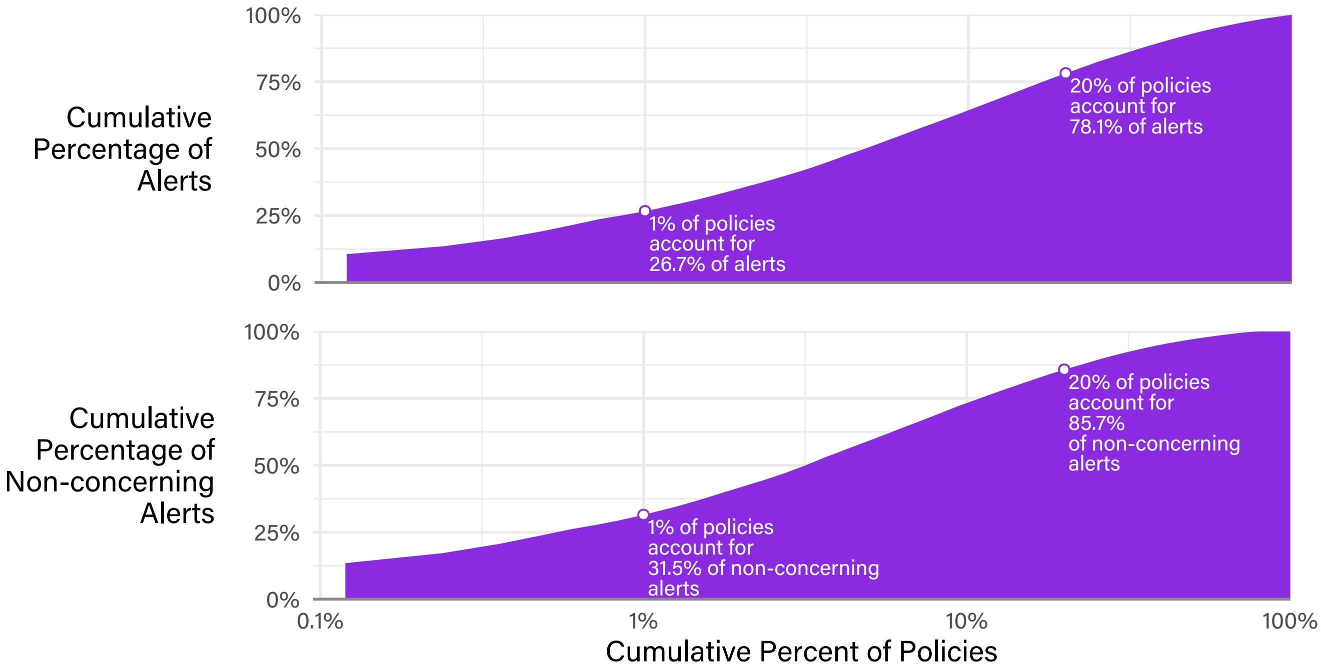

FIGURE 12 – PROPORTION OF ALERTS BY CUMULATIVE POLICIES

Figure 12 allows us to learn a lot. The law of the vital few. The 80/20 rule. Pareto, but for policies. There is a “heavy-tailed” effect with policies. Seventy-eight percent of what analysts actually look at comes from only about 20% of polices. Among the alerts they look at and then mark as “not concerning,” 1% of the policies account for 31% of those alerts, with just a fifth accounting for 86%. While this unbalance may seem dismaying at first, it’s actually great news. In particular, this means a small amount of tuning—e.g., not focusing on those policies prone to creating non-concerning alerts—can save a great deal of time and effort.

One obvious method for tuning is to simply only look at the “criticality” of the policy that generated the alert. After all, when creating policies, a subjective determination of the importance of the policy is likely made. Figure 13 breaks it down.

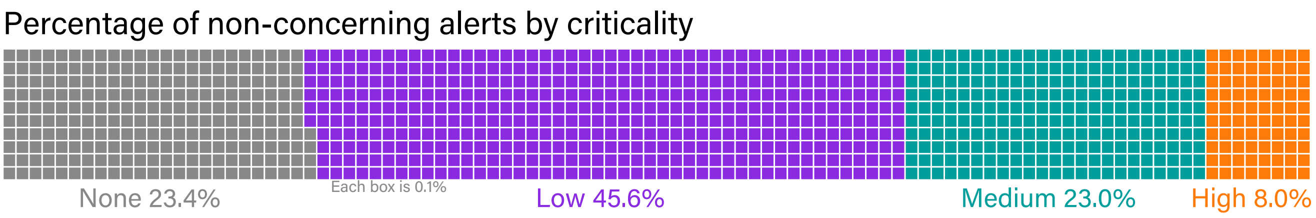

FIGURE 13 – NON-CONCERN RATE FOR POLICIES OF DIFFERENT CRITICALITY

“Low” criticality policies make up a large plurality of non-concerning alerts, with “High” criticality alerts making up a relatively small 8%. “None” and “Medium” are balanced with about a quarter each. This actually makes a good amount of sense that the sum of “Low” and “None” criticalities makes more than two thirds of violations marked as non-concern. “Low” criticality alerts seem like the likeliest candidates for “non-concern” as there is something there (it’s not “None”), but it’s not very important. Whereas “High” alerts should be relatively rare and should make orgs pay attention. No organization needs, or wants, to hear that the sky is falling so often that they don’t realize when the sky actually is falling, so it’s good to see that among non-concerning alerts “High” is the rarest.

One other note is that criticality may denote potential impact and not the likelihood that something bad is happening. An example: organizations probably want to show a "critical" alert every time the built-in administrator account from a Windows system is used. A good policy would be for this not to happen, and if it is an attack it means the attacker has obtained high privilege on the system. But, it turns out that, although against the policy, this is not an infrequent occurrence at many organizations. This could account for why some of those "High" criticality alerts are nevertheless non-concerning.

Let's look at this from the two other perspectives with which we've examined the data: the data source being monitored (Figure 14).

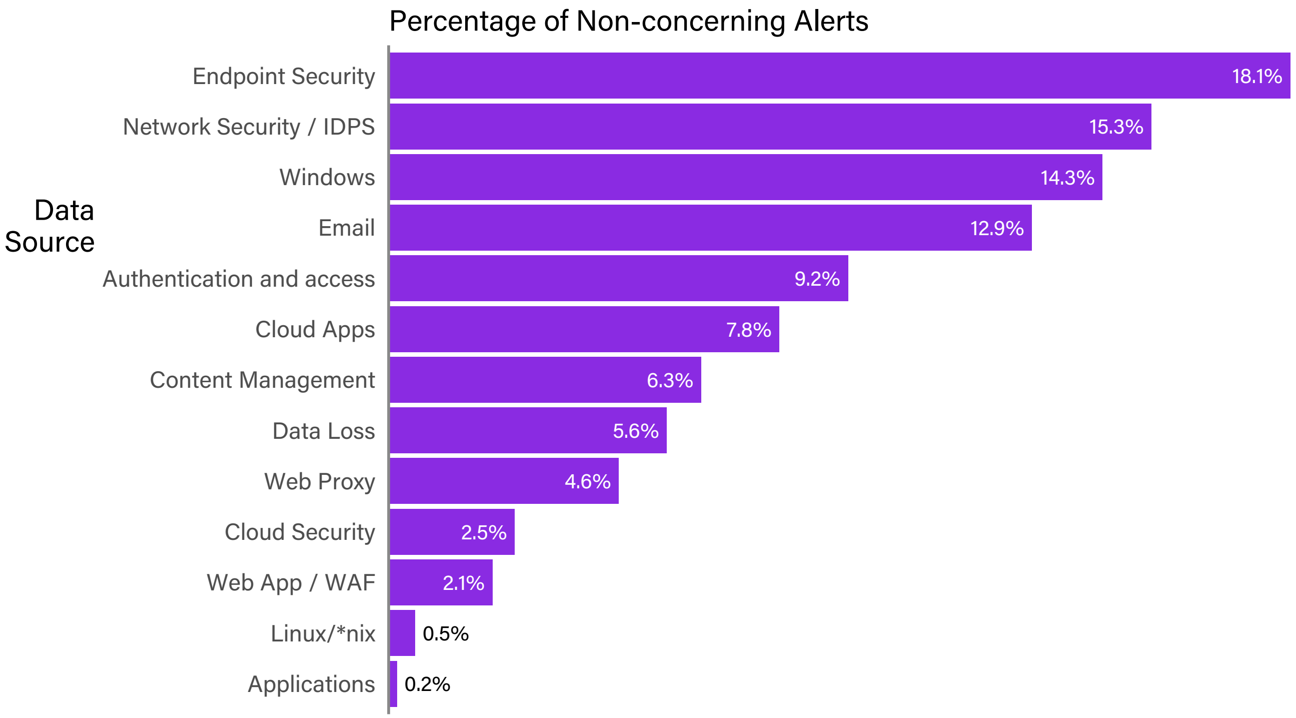

FIGURE 14 – NON-CONCERN RATE FOR DIFFERENT DATA SOURCES AND POLICY TYPES

At the top of the list are things that tend to focus on very common threat vectors, namely Endpoint Security and Network Security / IDPS. These types of data sources usually detect an enormous amount of what can generously be called "commodity threats", such as simple malware and botnet activity. Although most of these attacks are real, most organizations have security controls in place to simply block them, so it's unsurprising they make up more than a third of all non-concerning alerts. On the other end of the spectrum, "Applications" make up the lowest percentage of non-concerning alerts, at just 0.2%. These are specific applications flagging things that are definitely strange, like "abnormal number of records accessed" and "user adding privileges to self." Easy, fast rules that are likely to address violations that make an analyst worry.

APPLICATIONS
0.2%
NON-CONCERNING
ALERTS

So what are the "good" policies?

In order to identify a good policy, we have to first define what we mean when we say "good policy." If we take a look at the work we've gone through to get here, we can start to craft a working definition of a good policy. A good policy is a combination of two of the things we've seen already; it generates a lot of information and is judged by analysts as actually identifying interesting information on a consistent basis. We already have both of these measures: the violation rate and the non-concerning violations. Except now, we are going to turn that second one on its head and look at the percentage of alerts that a policy creates that are actually of concern. Take a look at Figure 15.

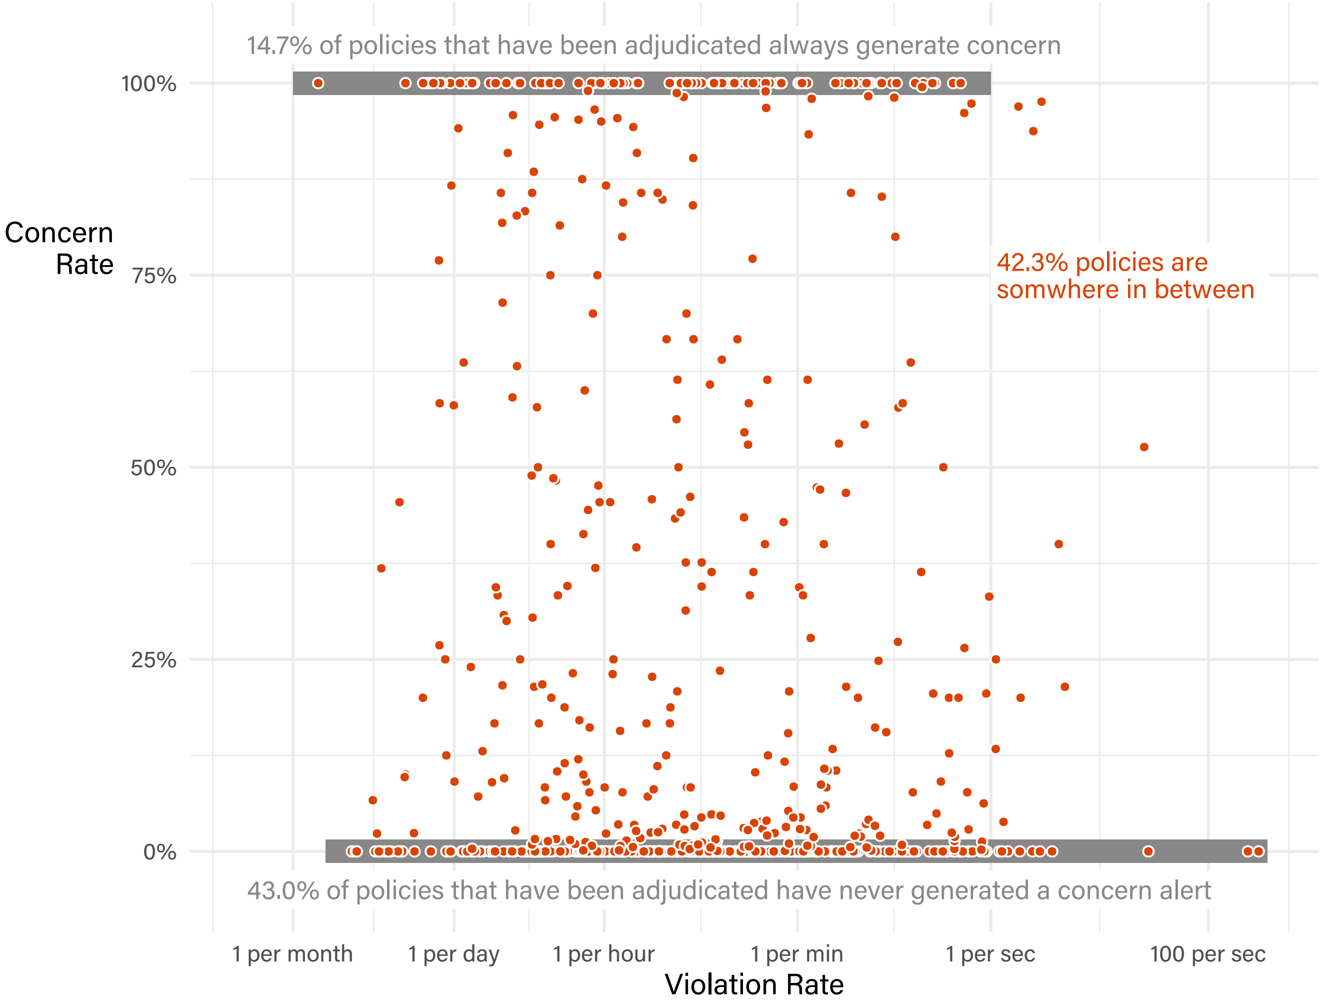

FIGURE 15 – CONCERN RATE'S (NON-EXISTENT) RELATIONSHIP WITH VIOLATION RATE BY POLICY

Figure 15 shows what us statisticians call "no correlation". That is, some policies are clearly very "chatty" but with a lot of good information (upper right), while others are mostly just babbling away at a mile a minute(lower right). On the other end of the violation rate spectrum, some policies are relatively quiet while generating few alerts (bottom left), while finally the policies in the upper left are generating a high percentage of concerning alerts with a parsimonious amount of violations. Given what we've seen before, namely data sources and policy types, what makes a good policy?

Features of good policies

First let's take all those points in Figure 14, categorize them by data source and examine how they are performing.

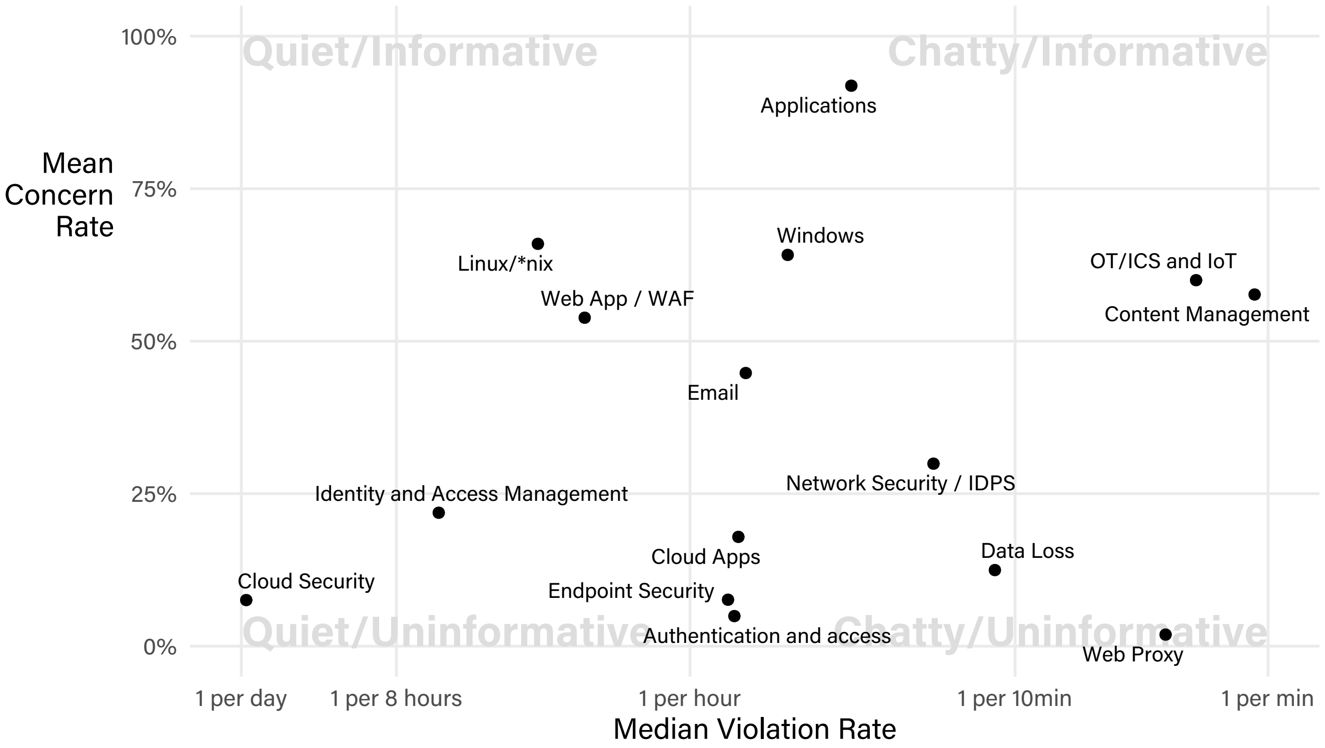

FIGURE 16 – CONCERN AND VIOLATION RATES FOR DIFFERENT DATA SOURCES

Figure 16 presents us with a better view of the state of the world. In the lower right of the figure, we have things that are extremely chatty, while not being terribly informative, particularly things like Web proxy and Data Loss monitoring. Similarly chatty is good old IOT and Content Management, but at least their concern rates fall above 50%. Cloud Security rarely generates violations and when it does, it's rarely of concern. It is interesting to note that there isn't really any data source that is generating a low number of high-quality violations.

Of course it may not be just what we are monitoring, but how we are monitoring. If we combine both the technique and the data source, we start to see a more complete picture.

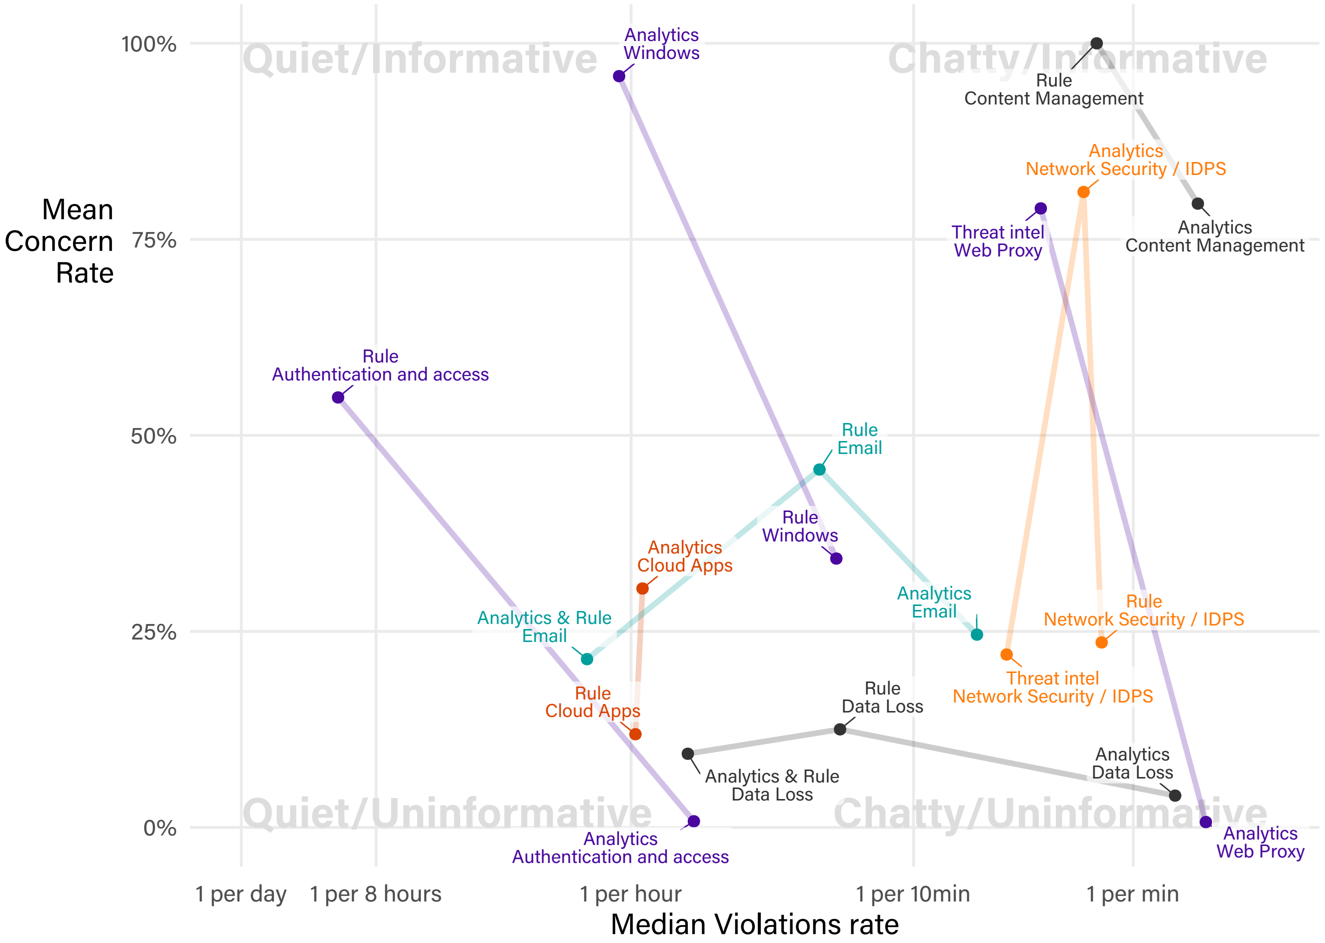

FIGURE 17 – CONCERN AND VIOLATION RATES FOR DIFFERENT DATA SOURCES AND POLICY TYPE

The points here are a little sparser, and we wanted to restrict ourselves to data source and technique combinations that had a minimum of five policies to calculate our statistics. Each color represents one of nine data sources, with a new point for a different technique used to monitor that source.

Some takeaways jump out immediately. For email (green), just using analytics can generate more than 16 alerts per hour, with most of them being non-concerning. Using rules, the fraction of alerts that are actually of concern goes up with a lower rate of violations. Combine them, and the violation rate plunges, but so does the concern rate. For a few data sources, some techniques are clearly better than others. For a Web Proxy, Threat intel is much better than Analytics. Rules are better than Analytics for Content Management and Authentication and Access. Conversely, Analytics excel for monitoring Windows, Network Security, and Cloud Applications.

Conclusion

We've taken a journey to get to this point right here. What's abundantly clear is that we know what we know, we don't know what we don't know, and the firehose of information can be overwhelming for many organizations to parse through. Threat detection is hard; understanding not only how to filter the information, but also what the information means is critical for success. What good is having an all-seeing eye if you don't know what you are looking for?

Understanding and measuring your own organization's systems is truly the first step in honing the efficiency and effectiveness of your threat detection. Once you are able to discern which organizational policies produce noisy with non-concerning alerts, you may find yourself having an easier time sifting through the alerts to find the concerning violations more efficiently. But even if you do that within your own organization, it's difficult to know if what you are seeing is specific to your organization or part of some wider trend. This is where cloud-based products that can see trends across multiple customers in multiple environments can provide insights into this problem that no one else can.

The next step? Taking a look at your organizational policies to see if there is a way to restructure or create new policies that can lead to more effective monitoring in the future. Partnering with a provider like Securonix can help even more, since Securonix can leverage data from their vast array of customers to assist you in crafting effective policies.

FINDING THE SIGNAL THROUGH THE NOISE

A COLLABORATIVE REPORT BETWEEN

securonix

119 Cyentia INSTITUTE