CHAPTER 3

Cost Benefit

In this chapter, we will explore the concept of cost benefit and how it applies within the cybersecurity industry. The term itself is rather self-descriptive. Cost benefit is an assessment of how beneficial something is when offset with cost. Typically, the benefit in cost benefit, when referred to in cybersecurity constructs, is the ability to mitigate risk. Risk could be risk of exploitation or other worries, but they all roll up to either financial risk or loss of life or both. Typically, the cost in cost benefit, when the term is referred to in cybersecurity circles, is a cost in dollars, but this does not necessarily have to be the case; sometimes, it could be in the form of resources or time spent, but ultimately, those too get rolled up into a dollar amount.

Good cybersecurity cost benefit is when you implement a cybersecurity product, capability, or subscribe to a cybersecurity service that mitigates enough cybersecurity induced that the cost of that asset is positively offset by the risk it mitigates. Bad cost benefit is when the cost of the asset far outstrips the potential benefit that could be gained from implementation. There are many defensible ways to metric cost and benefit and a dearth of calculative ways to determine whether a given cost benefit is good or bad. That is not the focus of cost benefit in this chapter or book. Instead, the takeaway is hopefully that you as the reader are able to assess cost benefit in your own terms, situationally, and as needed, but in a way that explores not just what a product, service, or capability aims to provide but whether in doing so it truly provides cost benefit at the strategic, organizational level.

J. G. Oakley et al., *Theoretical Cybersecurity*, https://doi.org/10.1007/978-1-4842-8300-4_3

Being able to appropriately determine cost benefit is the difference between being able to convince people or not that they should buy or use your cybersecurity thing. It is also the difference between procuring and implementing cybersecurity things that benefit your organization or not. Even in cybersecurity, it is all about sales, and selling in our case involves (or should anyway) proving true cost benefit of cybersecurity.

There is admittedly a slight difference in selling cost benefit between commercial and regulated spaces such as the DoD, HIPAA, or financial institutions. In commercial environments, you have to sell people (your boss, their boss, shareholders) on cybersecurity itself by using the best combination of products, capabilities, and services you can architect together. In regulated spaces, cybersecurity itself is required, so you are instead in the position of having to illustrate the cost benefit of a given product, service, or capability within the defined architecture the organization is already being held to such as the risk management framework (RMF) put out by the National Institute of Standards and Technology (NIST).

For the sake of brevity, I will start using the term cybersecurity thing, which is intended to represent a cybersecurity product, capability, or service. When we consider the cost benefit of a cybersecurity thing, we first need to ask five simple questions that may have very complex answers:

1. What is the intended specific technical cost benefit?
2. Does that specific cost benefit translate into organization-wide cybersecurity cost benefit?
3. Does that cybersecurity cost benefit translate into strategic cost benefit?
4. How long or at what point does it become cost beneficial?
5. How long does the cost benefit last?

Warning

I am about to make some statements that may been seen as contentious; you may not agree with them. If you do not agree with the following statement, I hope you read on and give me a chance to prove my point to you and outline exactly what I mean.

No one cares about cybersecurity.

There, I said it. I suppose we could clarify and say that specifically, cybersecurity vendors and cybersecurity consumers don't care about cybersecurity. The professionals on both sides of the producer consumer equation are often passionate about cybersecurity and enjoy the puzzling, problem-solving nature of the body of work we are a part of. Being blunt though, until we as cybersecurity professionals truly understand what really drives cybersecurity-related decisions by vendors and consumers, we will not be able to help them be secure despite themselves.

That statement is probably not going to make me a lot of friends among future cybersecurity employers, but let's get into the what and why of such an audacious statement.

Real Motivation

All right, so we have covered what is meant by cybersecurity cost benefit and that it is important when trying to offer or consume a cybersecurity thing. The next concept that needs to be understood before we talk about applying the good cost benefit analysis to cybersecurity things is a harsh reality, an ugly truth or whatever other label you want to give to what I just said.

No matter the lip service of any government institution or commercial organization, they really don't care about cybersecurity. They care about their own strategic goals, objectives, and outcomes. Sure, cybersecurity is seen as either a challenge to or a protector of those strategic outcomes, but

no organization actually has cybersecurity as a strategic outcome. If you are the Department of Defense, you are trying to save lives and protect the country, having good cybersecurity helps you make sure you can do that in contested environments like the cyber domain.

If you are a vendor like, you may provide cybersecurity to an organization like the Department of Defense, but your strategic outcome is to make money and continue your existence. Unlike those federally funded research and development centers (FFRDCs) or not-for-profit organizations like MITRE, you are still beholden to budgets to pay your employees and without focusing on that foremost the people running such organizations still risk them folding. This is not to say that vendor organizations' best path to achieving their strategic outcome isn't providing good cybersecurity for their consumers, I am just trying to get everyone to acknowledge that no organization has cybersecurity as a strategic outcome, and strategic outcomes drive cost benefit.

Examples

If you still disagree with me about my abhorrent statement about no one caring about cybersecurity I have some illustrative examples of this being the case that span the gamut of cybersecurity industry functional domains.

Industry Wide Example: Retention

This one is a pet peeve of mine, so I apologize for the tower of soap boxes we currently sit atop. It is my opinion that complaining about retention in the cybersecurity industry is obnoxious. I hear and see things like, "Well every time I train up a person or they get certified or finish their degree they move on to another position." Then you as the employer didn't try hard enough to retain them. I would think that if you took that person's new resume, and had it sent as an applicant to replace the person who you just let go, they would probably get about the same pay as that person

is getting at their new place of employment. Worse, you've lost the tribal knowledge that person has about the security apparatus they work within in addition to their technical skills and experience.

My main point though is not that I get annoyed about companies complaining about a retention problem they could solve themselves by promoting within and rewarding organic growth. The point to illustrate is that such companies (most if we are honest) show they don't care about cybersecurity because they let the tribal knowledge walk. It is worse, or maybe I am just more familiar with government contracting, but the story goes something like this:

1. Leverage the resumes of talented people to win a contract.
2. Hire as cheap of resources as possible to staff contract personnel requirements.
3. If they grow through certification or degrees of years of experience and want more salary, let them find other employment.
4. Hire the cheapest person possible to fill the same slot and bill the same rate to maximize profits.
5. It doesn't matter if they do well, the government is incentivized to pick a new contractor on re-compete anyways, so they don't look like they are playing favorites. They are also incentivized to spend all the money they set aside for that contract otherwise they can lose budget allocations for follow-on years. The work would have to be so poor that it became worth the government's time to kick the company off and re-allocate the funds, and that is almost never going to happen.

OK, that was a bit extreme, and maybe jaded. However, when cybersecurity vendors let contract positions become vacant instead of re-investing in their people, they are often doing so to maximize the profitability of that contract. Conversely, if they had the cybersecurity of their customer as a strategic outcome (which they don't, they are a business) they would ensure the tribal knowledge that would provide greater cybersecurity from such personnel stayed on the contract. It is not malicious or wrong for a cybersecurity vendor to have profitability as a strategic outcome and not cybersecurity. However, it is important as we evaluate where the cybersecurity industry is at for us to acknowledge this as a truth.

Defensive Cybersecurity Example: Metrics

If you have ever been a member of a security operations center (SOC) where detect functions are executed to provide cybersecurity, you may already know the point here. In many SOCs, it is more about cybersecurity theater than it is about providing actual cybersecurity. This is because in the best SOC, operated perfectly and run with the best tools, signatures, and by the best professionals in the most secure network, you would probably never get an alert.

I can tell you from personal experience it is very hard to continue to prove the cost benefit of nothing. So how do we try and communicate the value of our SOC to the people paying for it. The most common example I have seen is reporting metrics that sound impressive but have little to no cybersecurity meaning. Every reporting period, the SOC says it had some million number of events monitored, and they saw some hundreds of thousands of hits on their external firewall. Well, all those statements mean that their publicly accesible attack surface is constantly getting scanned and probed by countless Internet-based agents just like every other Internet IP and that they have a large network that produces lots of events. To an operations- or business-focused person though, that sounds like they are doing a lot of work.

Again, if we are honest, the professionals who set up, operate, and maintain the SOC may very well care about cybersecurity. The SOC provider cares about maintaining the SOC contract and the customer cares about checking a box that they have a SOC so they cover their butts. The cheaper the better as long as they can keep telling their boss that the SOC has millions of events covered and hundreds of thousands of firewall hits monitored. Even in organic settings where a company stands up its own SOC, the person who is in charge of the SOC personnel still wants to keep his or her team the same size or grow it, and wants to keep their job and insure their people's jobs.

Offensive Cybersecurity Example: Reporting

The last example I would like to bring up is the one I am most familiar with and which I speak about in my book Professional Red Teaming. There are countless times where an offensive cybersecurity event is carried out and the end results are ignored, thrown out, or destroyed. This is done for several reasons. The customer, if they are the head of IT or security, for example, may not have the funds to fix anything in the report and knows they won't get them even with the report as evidence. So, they have the assessment conducted so they can tell their boss they did it or check a compliance box and then they throw away the liability that is an offensive security assessment.

The example I use in my other book is, imagine a hospital gets an offensive cybersecurity assessment done and there are ten findings. Say they do have funds to fix everything, but it will take a three-year period to cover all ten remediations, so they prioritize them and get to work securing their network. Now, a little over a year in, they have remediated four of the ten findings and are working on the fifth when an attacker leverages finding six to get into their network and steal HIPAA data for their patients. One of the patients sues the hospital and subpoenas their security-related documents, which include the report from a year ago with the ten findings.

Imagine the optics in the court room when they say, "look you all knew about this vulnerability for over a year, and it was in a report you got from a security team and now it was used to compromise my clients' data."

Pretty bad optics, right? Probably a case the hospital loses I'd guess. It is situations like this that drive people to pay for such assessments so they can say they did their due diligence but often the findings are a liability for any number of reasons. Again, if we are honest, the ethical hackers may actually care about the cybersecurity of the hospital. The hospital itself cares about protecting its financial interests and the offensive security vendor cares about keeping its professionals employed and expanding its customer base. The great cybersecurity professionals are the ones that find ways to help make the hospital more secure within the constraints of neither the vendor they work for nor the hospital actually having cybersecurity as their strategic outcome.

Understanding Cost Benefit Perspectives

We have covered at a high level what cost benefit means and how the term applies within the cybersecurity industry. What I would like to do now is to show that within cybersecurity there are different ways of evaluating cost benefit depending on the perspective involved. It is essential to understand how cost benefit is evaluated by all those involved in cybersecurity to come up with truly appropriate evaluations of cost benefit.

Cost Benefit to the Target

The more familiar perspective for most of us when considering cost benefit is to do so as the target of potential attack. In this perspective, the focus is on the perceived value of various aspects of the organization and how much should be spent to burn down risk to those assets. In our case, we are talking about burning down cybersecurity-related risks to such assets.

Consider a credit bureau as our target. There are several major entities; I will not pick a specific one, so let's just call them TransExperiafax. Now, though TransExperiafax offers credit reporting, monitoring, and protection to the people whose credit files they keep, most of their money is made from data analytics based on the files they keep. Let's say Transexperiafax makes an annual revenue of 2.5 billion from selling their data analytics to other companies. Also, they were recently the victim of a cybersecurity breach, and when it was announced, they lost 10% of their stock market cap, which equated to a 1-billion-dollar loss for the year for their shareholders.

If we think about how TransExperiafax might evaluate cost benefit itself as a potential target of future attacks, those two values are probably key. The 2.5 billion annually and 1 billion due to a breach are likely to be the strategic cost benchmarks to determine how much they are willing to spend on cybersecurity efforts toward protecting those year-over-year values and mitigating or avoiding catastrophic events.

Using these numbers, maybe TransExperiafax decides they'll spend 1% of the 2.5 billion annual risk plus the 1 billion potential loss values, each year, spread over 12 months equally. Their cybersecurity spending would look like Figure 3-1.

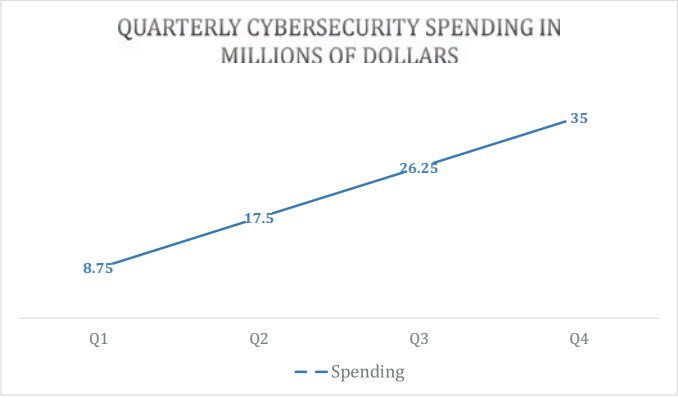

Figure 3-1. *Target Cost Benefit*

Cost Benefit to the Attacker

Unfortunately for TransExperiafax, the enemy gets a vote on cost benefit too. At least, they will have their own way of evaluating it. In what way does the attacker evaluate cost benefit? The easiest way for us to consider this is to make an assumption (probably a fair one) that the most likely malicious actor to target TransExperiafax is going to be an organized crime activity, potentially somehow tied to a foreign government, but not necessarily. If this is the case, then they are looking at TransExperiafax as a potential profit. TransExperiafax maintains some 500 million personal credit files and another 50 million company credit files. If we say the average company is ten people, that means there are essentially 1 billion personal credit files worth of data that they maintain. If the average credit file on the dark web sells for $5, that means the potential profit of compromising TransExperiafax is $5 billion. So, the attacker is going to evaluate the cost benefit of their malicious cyber pursuits against a potential $5 billion payout.

Using these numbers, maybe a criminal organization has decided they are willing to risk spending 1% of the potential 5-billion-dollar payout over a year, divided quarterly. Figure 3-2 shows what their cyber operations expenditure would look like

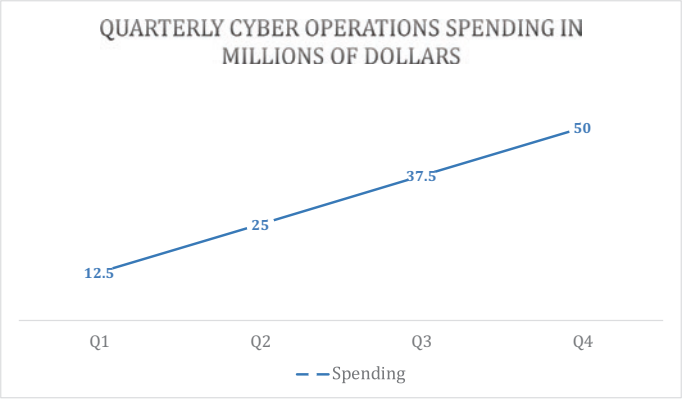

Figure 3-2. *Attacker Cost Benefit*

Summary

The point is that a target organization who is only able to see half of the cost benefit perspective picture is going to do cost benefit analysis on cybersecurity implementations without all the necessary information. If TransExperiafax did this, they would evaluate how much they should spend on cybersecurity using a 1–3.5-billion-dollar benchmark. Would they spend more or make different decisions if they knew that to the attacker, they looked like a $5 billion pay day?

Appropriate cost benefit analysis for cybersecurity products, capabilities, and services needs to at least consider both sides of this analysis and incorporate them into their decision process. Figure 3-3 illustrates the disparity in spending based on perspectives and shows that the attackers would always be spending more than the defenders. A key point too though is that this is just one attacker, maybe there are three, maybe there are many more, maybe they all go after TransExperafax this year, maybe each year, maybe consecutively.

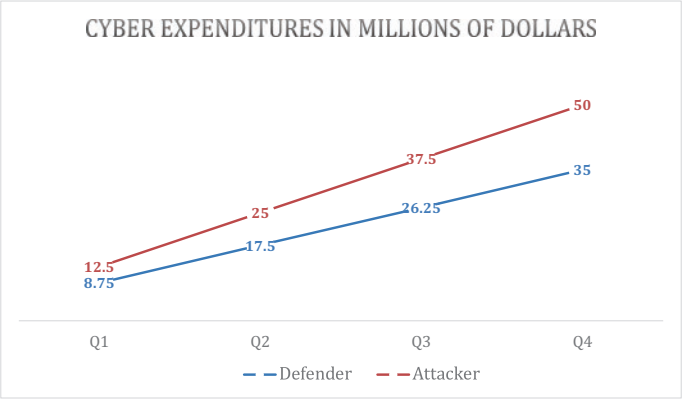

Figure 3-3. *Comparing Spending*

Understanding Cost Benefit Implications

While understanding the different perspectives of cost benefit is beneficial, it is also necessary to track the implications of implementing a given product, capability, or service. This means that even if on the surface a cybersecurity thing may look cost beneficial, we must also consider how the implementation of that thing alters the rest of our architecture. Primarily this means understanding how risk and work move around an organization as a result of such implementations and if those implications potentially negate the benefits of face value cost benefit analysis.

Risk and Work Are Never Destroyed (ish)

Much like matter, it is hard to destroy risk and effort. OK, with cybersecurity it is not so absolute. However, when cybersecurity things are marketed and sold and consumed, they are often done in a way that is dismissive of risk and work implications and instead focuses heavily, if not entirely, on the up-front change and cost benefit.

Poor Evaluation of Cost Benefit Implications

As an example of poor implications evaluation, let's consider the example of implementing automation through declarative languages. Put very simply, this is coming up with an easily understandable set of commands that, when executed by individuals, perform complex tasks behind the scenes. For instance, the declarative command 'newhost' might execute several scripts in the background to execute commonly performed tasks when a new virtual machine is stood up. It leverages a virtual machine API to create the virtual computer, it adds it to the domain, it creates a user profile on the machine, it installs antivirus and a suite of tools necessary for people to perform their job functions.

The face value of such an implementation is that instead of having a systems engineering team having to manually go through those tasks for each new machine they can simply type 'newhost' using the declarative language interpreter, and all the scripts and execution is orchestrated in the background through automation. This allows for the gaining of efficiencies in setup times for new machines, saves hourly wages paid to admins as they set up hosts and reduces the risk of mistakes by humans during the setup process that could make machines vulnerable. Easy sell right? Unfortunately, these benefits don't take into account where certain risk and work have moved to, and at what point the moves become worth it.

Good Cost Benefit Implications Evaluation

Using the same example, let's walk through what implications should be included in the cost benefit evaluation of a solution like this. There are several impacts to an organization when automation of this nature is put in place. Instead of spending a few hours here and there setting up a new machine for each new hire, an immense amount of work is put in up front to set up the automation mechanisms themselves. There is a clear need to understand what this work looks like and at what put the up-front cost starts to pay off and for how long it pays off. If it cost $100,000 in billable hours to set up the automation, but each machine set up only costs $100 in billable hours, we would have to have a thousand new machines set up before we start to see cost benefit in this regard. That might be easy in a very large organization; in a smaller one, it might not make sense to pursue this type of automation.

There is also the fact that such automation, scripting, and orchestration require specialized skillsets in a system engineering team that were not required before and additional personnel may need to get hired or time spent training them. These issues would add more impact to the cost benefit analysis on implementing this solution. There is also the

important movement of risk. Sure, human error is less likely to happen as humans are performing less actions. On the other hand, a single mistake at the orchestration level could not put every subsequent machine create at risk. There is another implication that must be considered in any cost benefit analysis: Where did the risk move to?

A Litmus Test for Cost Benefit

To this point, we have covered detailed methods and examples surrounding the concept of cybersecurity cost benefit, which in truth is simply true cost benefit for an organization. I would like to describe something I use as a quick litmus for cybersecurity things before I even go down the road of a comprehensive cost benefit analysis.

I refer to it as the 1-9-90 principle. The values may vary over time, but the point being made is that essentially, there are three types of threats that make up 1%, 9%, and 90% of cyber actors. Roughly 1% (probably less) of cyber actors are nation-state-level cybersecurity threats, another 9% are APTs and organized crime, and the other 90% are unorganized crime and script kiddies.

The 1% are undeterrable, unpreventable sources using almost completely if not completely unknown capabilities and the best way to deal with such risks are to find ways to accept that they could happen and find ways of living with them such as resilience and redundancy solutions.

The 9% are potentially detectable but unlikely to be preventable as they use both known and unknown capabilities.

The other 90% of cyber threats are those that must be prevented as they involve only known techniques and tools that can be scanned for and or caught by existing security tools.

So how does the litmus work? Well, if someone says they are developing a tool that can prevent nation-state-level APTs or detect them, you should take that claim with a grain of salt. As a purchaser or implementer of such a technology, you risk having sunk cost and resources into something that can't possibly deliver on what it claims. This 1-9-90 principle can be a great guiding resource for R&D as well, as you should focus on developing solutions that are aimed at mitigating specific threat actor sophistications in the most efficient and feasible ways. On the other end of the spectrum would be someone saying we should just accept the risk of the 90%; when you could easily thwart such known capabilities, why would you spend money on being resilient against them? Figure 3-4 illustrates this principle through a simple matrix.

	Tools / Techniques	Type of Actor	Risk Mitigation
1%	Unknown to public	Nation state	Accept
9%	Unknown and known	Organized crime	Detect
90%	Known to public	Unorganized crime	Prevent

Figure 3-4. *1-9-90 Principle*

If we look back at the moving target defense (MTD) example from Chapter 1 and applied this litmus, we probably wouldn't have to bother with further analysis. As the concept claims to PREVENT a 1% capability like a zero day, it would fail out litmus as striving toward an inappropriate method for risk mitigation. Of course, although 1-9-90 could be .001%, 9.999%, and 90% or have some other variance, the point is more that the majority of threats can be prevented; we should try to make sure we detect those that can't be prevented, but we should also acknowledge that there are unpreventable threats that we need to find a way to accept by being resilient to their manifestation. As with any rule or principle, there are surely exceptions; this is simply a quick sniff test ability to provide litmus to the cost benefit analysis of a given cybersecurity paradigm.

Summary

Turns out business thinking or operational thinking is really necessary to understand cybersecurity cost benefit. Additionally, we need to understand the cost benefit of a cybersecurity thing from the perspective of the defender and the attacker in any scenario. We also need to make sure we follow through on in-depth analysis of secondary and tertiary impacts on the resources and risk of an organization after a new cybersecurity thing is implemented. The 1-9-90 principle can enable an efficient, quick litmus to cybersecurity things and their potential cost benefit. Trying to be more secure is not always the right answer, and face value gains in efficiencies or decreases in risk are not always the full story. In later chapters, we will discuss some theoretical cybersecurity concepts that aim to provide real cost benefit. Even though, at the end of the day, our industry, like any other, is about business and not about cybersecurity at all, that doesn't mean the body of work itself can't be. Further, as an industry, if we can do a better job of putting forth feasible cybersecurity things with true cost benefit by leveraging the right kinds of people and context in our theoretical work, cybersecurity vendors, consumers, and professionals will all be the better for it.