CHAPTER 1

Introduction

The motivation behind writing this book was the observed lack of theoretical cybersecurity in the field and the resulting gap in innovation, security, and cost-benefit. I have had the privilege several times in my maturing career within cybersecurity to be afforded the opportunity to serve in job functions where my role was largely to assess, evaluate, and postulate improvements to cybersecurity frameworks that mitigate risk for large organizations. This included both commercial Fortune 500 companies as well as government organizations with global enterprise and loss-of-life and mission-critical systems.

I often found myself talking with other senior cybersecurity professionals, including the other authors of this book, about why the field seems to lack good theoretics. In the following pages, I will attempt to do my best to outline and describe exactly what is meant by theoretical cybersecurity. Next, we will identify some foundational issues of the field that have led us to what is being described as the theoretical gap. Lastly, we will posit some thought experiments and theoretical concepts of our own.

What Is It?

Theoretical cybersecurity is the proposed branch of cybersecurity where abstractions of actual technologies, systems, and organizations are used to rationalize, explain, and innovate upon the body of work that is cybersecurity. Specifically, I mean toward the improvement of the trade, of the craft itself. As an analogy, consider early automobile manufacturing.

J. G. Oakley et al., *Theoretical Cybersecurity*, https://doi.org/10.1007/978-1-4842-8300-4_1

The craft would be the designing and producing of automobiles. The technologies would be things such as riveters, welders, and paint guns. Improving the craft of automobile production involved the establishment of assembly lines. This is the type of improvement that I consider to be related to the craft rather than the technologies. I think in cybersecurity, we think we are doing a good enough job at innovation because we are constantly trying to one-up old technologies. Our field and those we protect would be the better for it if we also bettered and more often improved our craft itself.

It is also worth noting that good theoretical work is not just unadulterated and unbound pontification. Theoretical exploration of a craft should be rooted in the realistic constraints of the craft itself. This is not to limit the creativity of those doing the thinking. It is instead to ensure that when a theory is matured and ultimately experimented upon, the experiment is able to be designed in such a way as to prove out that the theory should be accepted as an improvement to the craft. Theoretical physicists, for instance, may think great and creative thoughts about the way things work at the macro or quantum levels. Their theories though must still be tied to the reality of their craft, in this case mathematics, such that when they conduct experiments on the theories, they are feasible, defensible as scientific proof. Otherwise, a theory could be posited, and an experiment created to prove or disprove it, but people may not necessarily believe the results to be true one way or another because there is too much detachment from reality. Once a theory has been proven, it can be taken by the applied physicists and become part of regular applications in the real world. Taking this back to a cybersecurity point, we need good theoretical cybersecurity that is both innovative as well as feasible enough to be proved out and make it to application across the broader field, or it is likely to be fruitless. Figure 1-1 shows out this process for theoretical cybersecurity, which is nothing different than many other scientific fields.

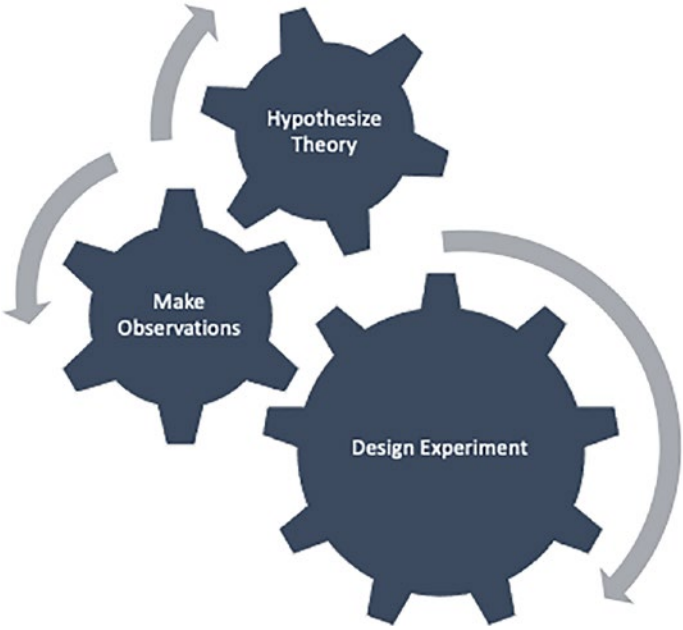

Figure 1-1. *Theoretical Cybersecurity Process*

Figure 1-2 shows the relationship between the theorization of a cybersecurity concept and its ultimate implications into the applied field.

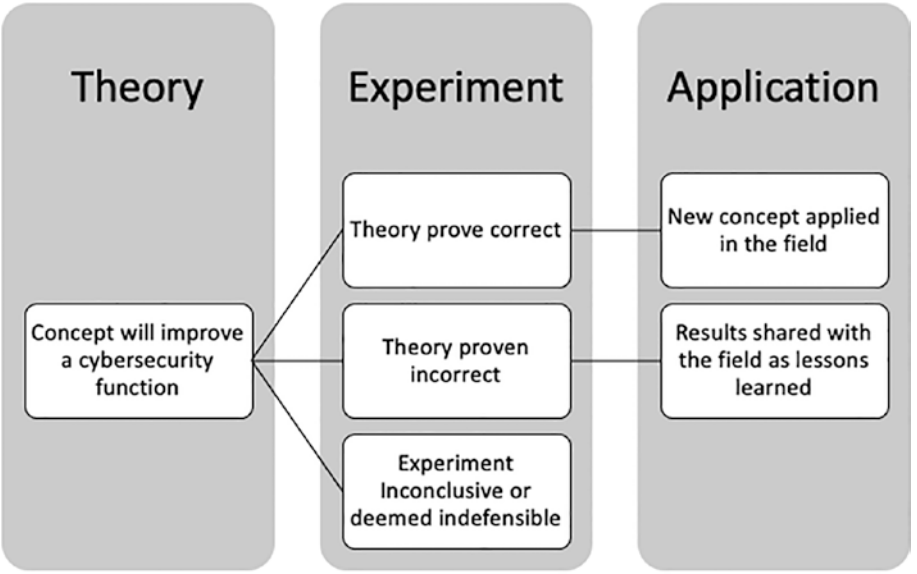

Figure 1-2. *Theory, Experiment, and Application Visualization*

It is vital that thought and experimentation is not wasted even when a concept is proven to be incorrect or is yet unproven. Disseminating knowledge of why a new theory was proven incorrect can still help inform the decision process of those professionals working regularly to apply known cybersecurity concepts. If an experiment was unable to prove a theory, this can act as a feedback loop for the theory itself and lend in refining the theory to a more provable state that is tied to reality.

What Is It Not?

Though the rest of the chapters will focus more on what theoretical cybersecurity can look like, why there isn't more of it, and so on, I think it is extremely useful to go through a case study of what it specifically is not. This is largely going to explore how important it is that the observation involved in leading to theoretical exploration of cybersecurity trade

craft be informed by science and tempered by experience in the craft. In cybersecurity, the underlying science may be things such as mathematics or hard computer science. Unfortunately, when those sciences stray out of technological improvements and into trade craft discussions aimed at improving cybersecurity, the lack of experience and context can lead to a lot of wasted time and moot ideas.

Case Study

In general, theoretical work in any field is the application of the scientific method to support or reject a hypothesis. Though a bit of an oversimplification, this is true as well in cybersecurity. The foundational issue that I will illustrate next is that the scientific method must be applied by people who are not only capable of rigor and academic thought in the theorization and experimentation but also that are informed and experienced journeymen or masters of their craft. For instance, a mathematician with a bachelor's, master's, and PhD in math would certainly be considered a journeyman or master of the science of mathematics. In this case her craft is a science. On the other hand, someone with a bachelor's, master's, and PhD in cybersecurity, but no real-world experience, would assuredly not be considered a journeyman or master of her craft. In this example, the craft is cybersecurity, which requires a certain understanding of computer science and other scientific concepts to perform, but which is itself a craft rather than a science. This is in part due to some fundamental flaws in the way academia has approached cybersecurity as a cash cow more than a defensible pursuit, but more on that later.

More often than not, the folks attempting to theorize on cybersecurity concepts are academics with experience in computer science and other fields, but without a journeyman-level grasp on the body of work that is the field of cybersecurity. This leads to examples like the one I am about to walk through, and it is important to understand how the following differs

from what will be prescribed later. The case study involved is indeed a theorization and an experimentation, but we must evaluate it through the lens of the craft and not just the technologies or, as you will see, we risk heavily investing in missteps. The following are not tied to any specific concept that has come forth but are close to several I have assessed in recent years. The case study does well to detail how attempts at innovation and theoretical cybersecurity can go awry.

Observation

Zero-day exploits give attackers a leg up on defenders because it allows them to come at defensive targets from previously unknown vectors. This makes it hard to be prepared for rapid pivoting, such as may be done by worms with a zero-day contained in them.

Theoretical Concept

To mitigate the impact of proliferated pivoting across an organization that is subject to an attacker leveraging a zero-day remote code execution vulnerability, systems should change constantly to make it harder for an attacker or an automated work to pivot around as the environment that is being attacked will constantly change. A changing attack surface to thwart malicious activity is known as a moving target defense or MTD. A way this could be done would be through randomly altering the ports used by certain services so that when an external entity throws an exploit against a port it sees, it is unlikely to succeed because the port is mapped to a different port on the host itself. Or the exploit may never get thrown because the attacker doesn't see the vulnerable service running because it is on an atypical port. Figure 1-3 shows how even in a linear network, where hosts can only be targeted after the attacker has gained access to the connected device, a zero-day-using worm can allow unfettered access.

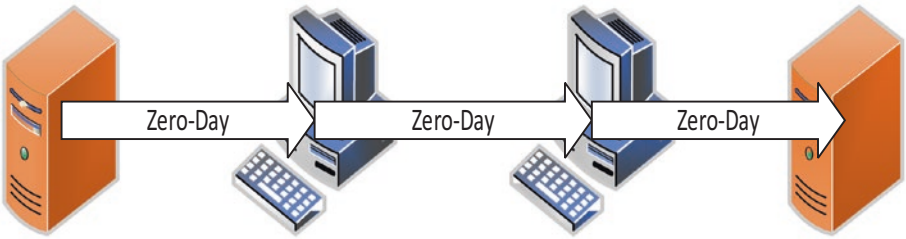

Figure 1-3. *Worm with Zero-Day*

This is representative of how the zero-day MS17-010 which targeted the Server Message Block version one protocol (SMBv1) and could have been used by a worm, such as it famously was in the Wannacry campaign, to exploit and pivot across machines running the vulnerable service on TCP port 445. Figure 1-4 shows this refined depiction of how a worm using MS17-010, which started as a completely unknown vulnerability, could pivot without pause since machines running SMBv1 on TCP port 445 had no idea they were vulnerable

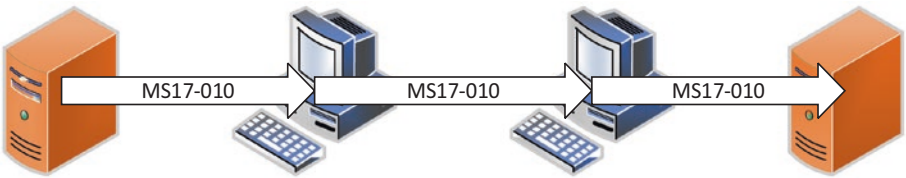

Figure 1-4. *Worm with MS17-010*

Experiment

In our experiment, we will use a Network Interface Card (NIC) adapter that has been programmed to take ports listening on a local host and represent them to the rest of the network in a randomized fashion. This way to users on a given host, their processes are opening the expected ports. So, running a command to show what ports are listening on any of the below servers would show SMBv1 on TCP port 445, but if it were scanned on

port 445 from another host on a network, TCP 445 would look closed. The NIC adapters all communicate to each other using an encrypted channel to share what the new random ports are and the NIC adapter seamlessly alters communications in transit. A network capture on the network connecting any two of these devices would now show traffic on TCP port 445 either. This NIC adapter to my knowledge does not exist, but it does well to give us a method for employing the MDT in our experiment, bear with me. Figure 1-5 shows how the ports for SMB would be represented on the network.

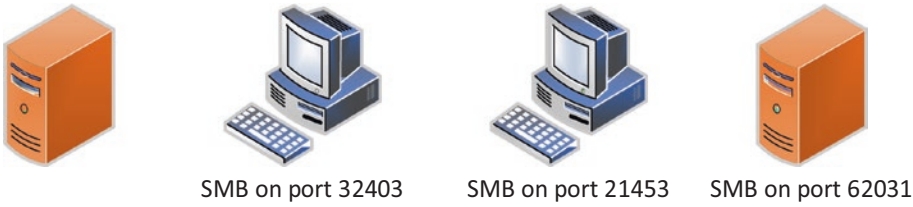

Figure 1-5. *Randomized SMBv1 TCP Ports*

The experiment will involve having a scripted worm that tries to throw MS17-010 in across the linear network and attempt to pivot and re-throw the exploit to delve deeper into the network. So that the worm can act as if it is still leveraging a yet unknown zero-day, the machines involved will be running SMBv1 on port TCP 445 that is vulnerable to MS17-010. This would be run once without the NIC adapters as a control, and the results looked exactly like Figure 1-3. The worm exploited its way all the way to the server deepest in the network by throwing the SMBv1 exploit on port TCP 445 after scanning and finding it open. Next, the exploit is run against the same network, but with the MDT NIC adapter turned on, it presents an attack surface as shown in Figure 1-5 to the worm.

Results

Because the worm scanned to see if TCP 445 was open and then would throw the exploit at the target it was unable to pivot to a single machine. We can try to limit the worm less and have it thrown the exploit without scanning. In this case, it was thrown once from the attacker machine, but failed to gain access to the next machine because the worm threw MS17-010 on its known port of TCP 445, but the NIC Adapter presented SMBv1 to the network on TCP port 32403. The results are shown in Figure 1-6.

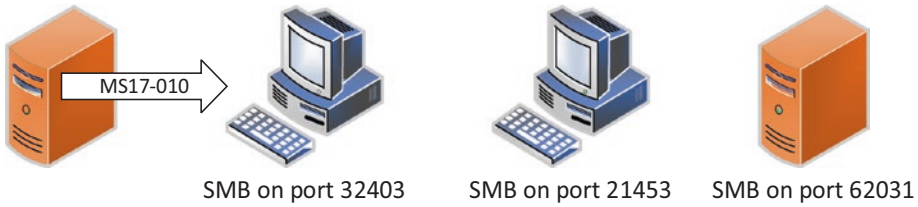

Figure 1-6. *NIC Adapter Run 1*

Taking it a step further, we then make the network non-segmented to see if our MDT NIC Adapter is capable of providing the same protections in a network not also protected by segmentation. Figure 1-7 shows that even in this situation, the worm threw the MS17-010 zero-day exploit three times unsuccessfully because the SMBv1 protocol which was vulnerable is running on unexpected ports.

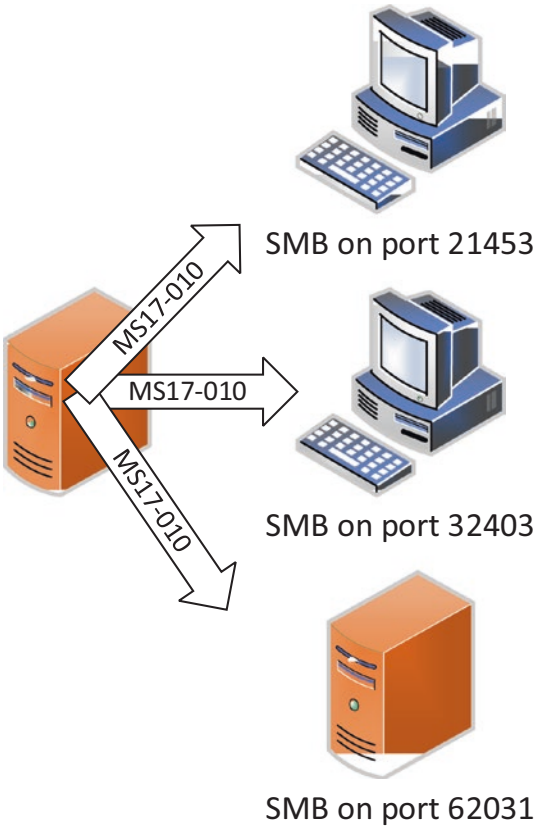

Figure 1-7. *NIC Adapter Run 2*

Conclusions

Based on the results of this experiment, it is logical to conclude that the hypothesis that the cybersecurity technology concept involved in the moving target defense NIC Adapter was successful in stopping a worm that leverages a zero-day. This experiment can be scaled and repeated and provide the same results, traits that are a hallmark of good academic experimentation.

This experiment exercised the scientific method on a technology that implemented the conceptual theory of moving target defense in an example network. So, what is wrong with it? I will walk through dissecting this shortly, but at a very high level, this theory and experiment is computer science oriented and not cybersecurity oriented. I have no issue with this experiment or its results. I do see an issue with it being painted with the broad brush that has become cybersecurity. We get into dangerous waters when we use non-cybersecurity experiments and outcomes to make cybersecurity claims.

The title to the paper or grant application that might come as a result of this successful experimentation should probably be something as follows:

> *Preventing Automated Communication with TCP Port Randomization*

What I think it would actually be called to garner attention, funding, and aim at productization:

> *Preventing Zero-Day Attacks by APTs with Moving Target Defense*

The difference here is pretty clear, the latter title is definitely written with a cybersecurity lens applied. I think it is easy to say that using this experiment and the language of a title like that could allow our MDT NIC Adapter to be shown at a large cybersecurity conference as a product offered by a vendor. They would even run the simulation on a loop showing how the MDT NIC Adapter prevented an automated attack with a simulated zero-day over and over. Since the MDT NIC Adapter does not affect the host machine it connects to, there is no integration cost on a per host basis. Just unbox, plug in between the Ethernet and the machine, and sync the adapters across your network and Bam! You are protected from APTs and their worms with the academically proven protection of our MDT NIC Adapter.

Case Study Analysis

As was already mentioned, there is a foundational issue in this case study of using observation and experimentation that lacks the perspective of the actual craft of cybersecurity to make statements about how to better the craft through something like moving target defense. Now we will analyze the case study to detail the implications.

Cyber Sniff Test

Let's take a quick look at this whole case study through the craft side of the lens instead of purely scientific.

Observation and Theory

In our case study, the scientists made an observation that zero-days are problematic because the attack surface they target is not known ahead of time. This means that it is extremely difficult to prepare for. The novel thought that was had by the scientists was to hamper an unknown dynamic threat by presenting it with a dynamic and unknown attack surface. The theory is then that a dynamic attack surface via random port mapping with our MDT NIC Adapter on a networks host will slow down, if not wholly prevent, zero-day attacks and those that use them from pivoting around a network.

Nothing in the previous paragraph is factually inaccurate. Where it errs is a lack of context to inform the scientists on the reality of the problem they are theorizing at fixing. Let's walk through the actual type of threat this theory would be aimed at mitigating.

- It wouldn't stop external initial access such as would be established by phishing and other malware campaigns.
- It wouldn't stop an insider threat.
    - They would know of the technology and circumvent it.
    - If the MDT NIC was seamless, attacks launched from an insider computer would communicate without issue.
- It wouldn't stop an advanced persistent threat (APT) or even a less sophisticated human with interactive access.
    - They might have the context of an insider threat and those two points apply.
    - They would scan all ports and see that something odd was going on and work to circumvent it easily.
        - SMBv1 would still answer normally regardless of the TCP port it was communicating on, otherwise it couldn't function.
        - OS/Software fingerprinting available in free open source scans would likely catch this right off the bat.
- It won't stop an internet-based worm OUTSIDE the network using a zero-day because the remote hosts accessing internet facing resources won't all have access to something like the MDT NIC and therefore it can't run on internet facing attack surface or it would be problematic operationally.

- It might stop an automated attack such as a worm INSIDE the network with poor automation logic as it would fail as described in the experiment.
  - Though it is likely a tool deployed with a real zero-day was done by an interactive APT hacker who would notice its failed proliferation and adjust accordingly to circumvent it as mentioned earlier.
- It does not address less sophisticated actors that do not have zero-days for the same reasons.

Table 1-1. *Threats Prevented*

Origin	Vector	Prevented?
External	Phishing campaign	No
External	Worm using zero-day	No
External	APT using zero-day	No
External	Hacker using known exploit	No
Internal	Insider threat	No
Internal	APT using zero-day	No
Internal	Hacker using known exploit	No
Internal	Worm using known exploit	If it uses poor logic
Internal	Worm using zero-day	If it uses poor logic

In my opinion, the only realistic threat this would stop is an automated worm with bad automation logic, as is shown in Table 1-1. Zero-day exploits are extremely expensive, rare, and the first time they are used they are not nearly as effective, as the cat is likely to be out of the bag. This doesn't sound like the kind of resource that would be blindly deployed at all by sophisticated and nation state malicious actors, let alone via a worm with poor exploit logic.

Even if you disagree with some of my assumptions or conclusions, I think we can agree that in general, this observation and theory lack necessary context to be rooted in reality despite the supporting science. Further, the solution developed doesn't really have an existing problem to solve if you understand the way actual threats operate. This is a clear illustration of how science without cybersecurity context can lead to observations and theories that even if proved out such as ours was, don't lead to any meaningful contribution to the field.

Experimentation

Any experimentation to prove out theories and observations made from an inadequate perspective are bound to yield results that end up being of little value to the craft of cybersecurity. As we have just seen, they may prove out a theory and result in the successful marketing and dissemination of a tool or technology, but the likelihood that such a tool would actually mitigate cybersecurity risk are low or coincidental. Later in this book, we will spend an entire chapter on designing good cybersecurity experiments, but first we will cover what it will take for the field to begin generating more theoretical cybersecurity paradigms.

Implications for Implementation

One last thing I will cover in a small amount here and which will be covered later is how there is a distinct need for understanding implications and true cost benefit. Whether observation, theory, and experimentation are done purely scientifically or with context of the craft, implications of advertised cost benefit are not often sufficiently incorporated into evaluating true cost benefit. Even if we thought our MDT NIC Adapter provided improved mitigation, we need to take it many steps further in understanding the cost benefit of such a technology.

One quick set of examples would be the downstream implications to other cybersecurity apparatus post implementation. Our MDT NIC adapter randomizes exposed TCP ports, and the traffic will travel to the new random ports between hosts. This means that any network-based intrusion detection system (IDS) or security information and event management software (SEIM) won't be able to leverage their heuristics or other capabilities because the traffic will look odd to them too. Further, actions such as forensics activities and threat hunting will also be hampered by having to tie what was captured or seen on the network with logs of what was the state of the MDT randomization at the time of other events.

With just these quickly mentioned implications to putting in place even an MDT NIC Adapter that did provide actual security mitigation, is it likely to do so to the extent that it is worth addressing the other sunk costs and movement of risk and work across the attack surface? I think not. This is not always the answer, but this is the bare minimum extent to which implications must be taken into consideration, even for valid theories and concepts, when we evaluate the cost benefit of a novel cybersecurity idea. In a later chapter, we will deep dive on understanding and evaluating true cost benefit in cybersecurity. This is an integral part of the craft, this is the way.

Summary

We have touched on the concept of theoretical cybersecurity and walked through a detailed case study to really hit home what is meant. The typical academic and industry scientific theorization and experimentation are not what I would call theoretical cybersecurity, they are more exploration of a particular science involved in information technology. We have discussed that what is most important in cybersecurity in general and its theoretical endeavors specifically is that it is a field which requires a strong scientific

understanding of various technologies and concepts as a starting point, but which should be measured largely based on experience in the craft. The case study was used to show the separation between science and craft from a theorization perspective and how it leads to consequential impacts to the field in general.

It could be argued that if theoretical cybersecurity were only performed with a journeyman type of perspective, many potential academic achievements and innovations may be missed as thus limits the researcher population. Perhaps, the preceding case study leads to some other finding that ends up benefiting the field. That is more than fine, and I hope it does. What I am shouting from a soap box to the industry is that such an example is not theoretical cybersecurity and that we need more of what I have deemed theoretical cybersecurity if we are going to achieve meaningful improvements to our craft that aren't. As we will see in the coming chapters, it is hard to get the right people to make the observations and come up with theories, and it is hard to come up with good cybersecurity experimentation because when it is done right, it requires science and a lot of human involvement, which makes defensibility challenging. Most importantly, much of what is yet to be covered focuses on understanding, providing, and ensuring cost benefit as an outcome of cybersecurity, which should always be evaluated through the amount of cybersecurity risk mitigation necessary to secure strategic objectives of any organization.