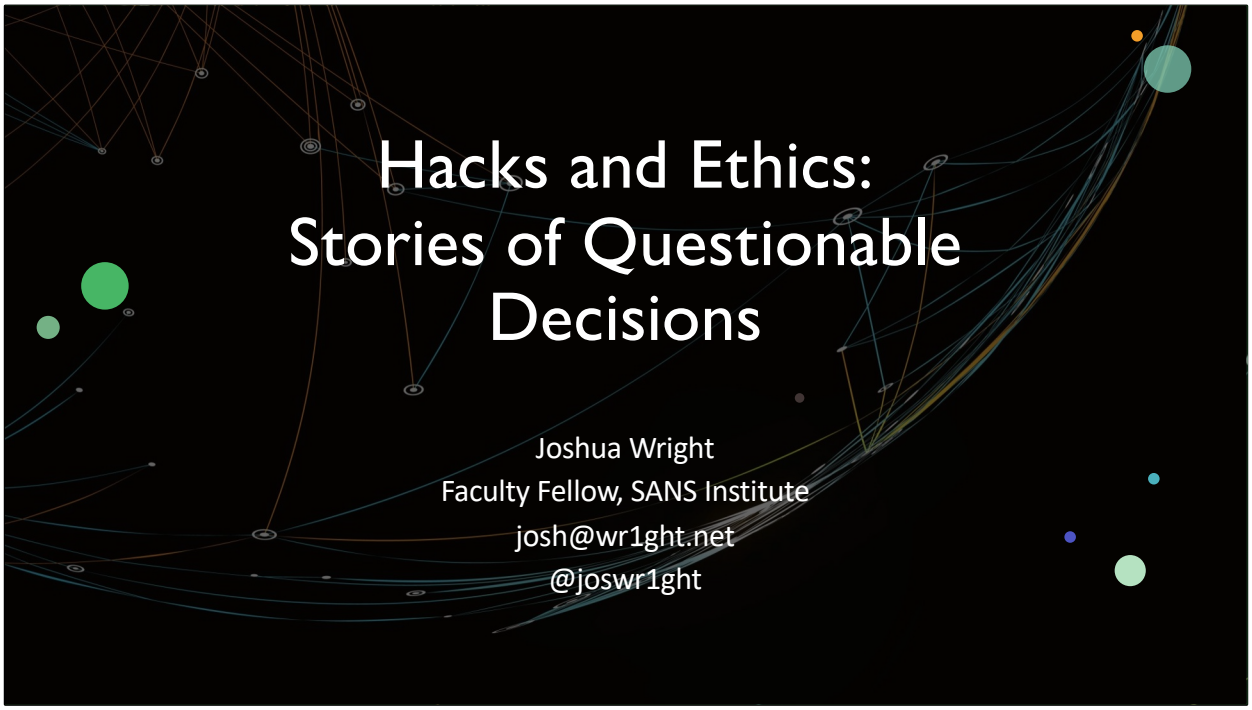

In this presentation I'll look at three ethical challenges I've experienced with my career, peppered with stories of hacks. I'll introduce several frameworks and techniques for the judgement of ethical behavior for each, pointing out the challenges I've experienced.

Joshua Wright
Faculty Fellow, SANS Institute
josh@wr1ght.net
@joswr1ght

Many years ago, I received a phone call from an unknown caller. This was in the days when we still answered those calls. Two people were on the phone when I answered. Lawyers from Cisco Systems.

I had recently returned from speaking at my first DEF CON hacker conference in Las Vegas where I was critical about a vulnerability I discovered in the Cisco LEAP wireless authentication protocol. This was at a time where Wi-Fi security was still a mess, and Cisco's proprietary solution was a key product differentiator that bolstered their product line. My analysis showed its faults -- in many cases, Cisco's security was worse than other widely-available protocols. I was preparing to release a straightforward exploit tool.

At first, I thought *maybe they are going to offer me a job*, but that's not what happened. "Cease and desist", "investigating a lien on your house", and "pursuing criminal charges" are snippets of what I recall from the conversation.

At the time, I was a new dad, and my wife and I were fixing up our first home together. I worked for a small university in the IT department. These actions by Cisco were threatening to my family's well-being.

Sort of how my releasing an exploit tool to take advantage of an un-patchable vulnerability in a widely-deployed wireless technology was harmful to many others. It forced me to start questioning the importance of ethics in cyber security work.

Image source: https://youtu.be/IWYb_gWCg1k?t=142

In the field of cyber security, there's a lot of good reason to think critically about our actions, to ask questions about ethics and the impact our decisions will have.

- We have more access to sensitive and company data than many other people
- With that access comes great responsibility
- Behavior that is unethical isn't always easy to identify; it's easy to make decisions that are unethical without knowing
- As we come to rely on increasingly complex technology, I believe we will see new and more complicated challenges in applying ethical behavior

In this presentation I want to share three case studies of times where I was personally challenged when making ethical decisions, to look at some strategies we can employ to evaluate our ethics, and some insight on how we can improve our own ethical decision making.

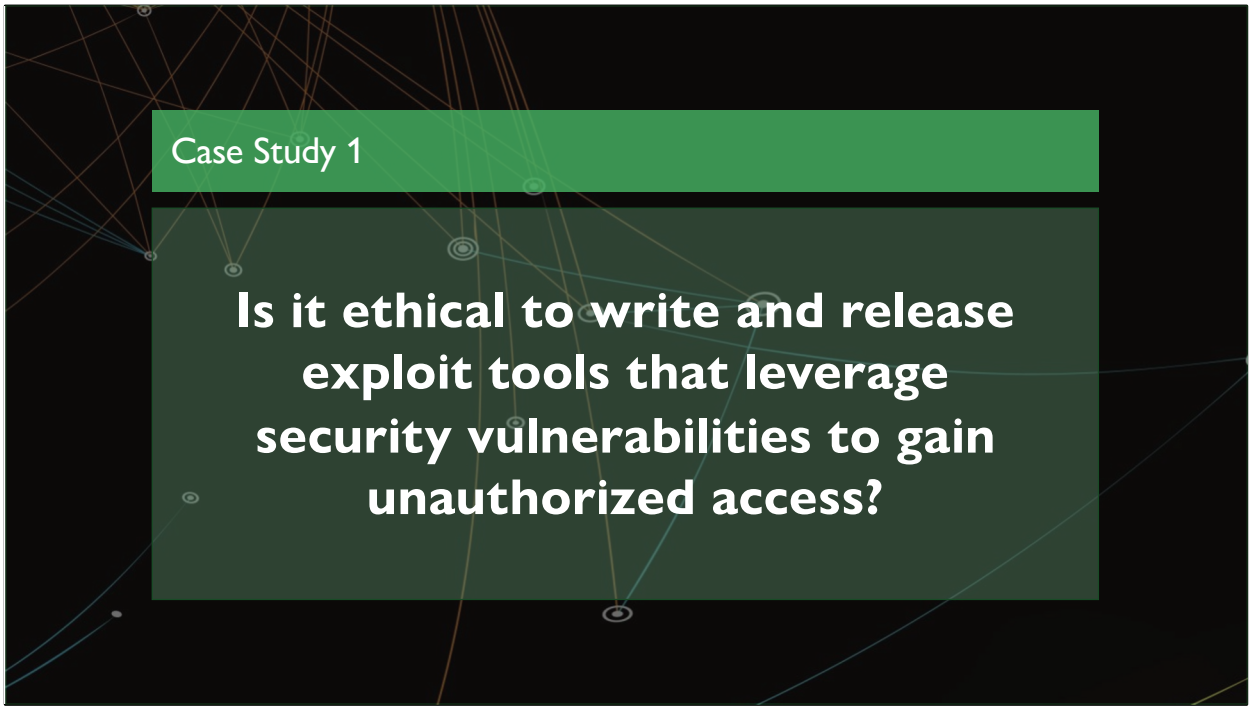

Case Study 1: Ethics as an Offensive Security Tool Author

Is it ethical to write and release exploit tools that leverage security vulnerabilities to gain unauthorized access?

Exploit tools can be used by people with authorization to test the security of their organization, but they also can be used by people with ill intent to defraud organizations or to steal from individuals.

We can apply ethical frameworks to help us in answering this question.

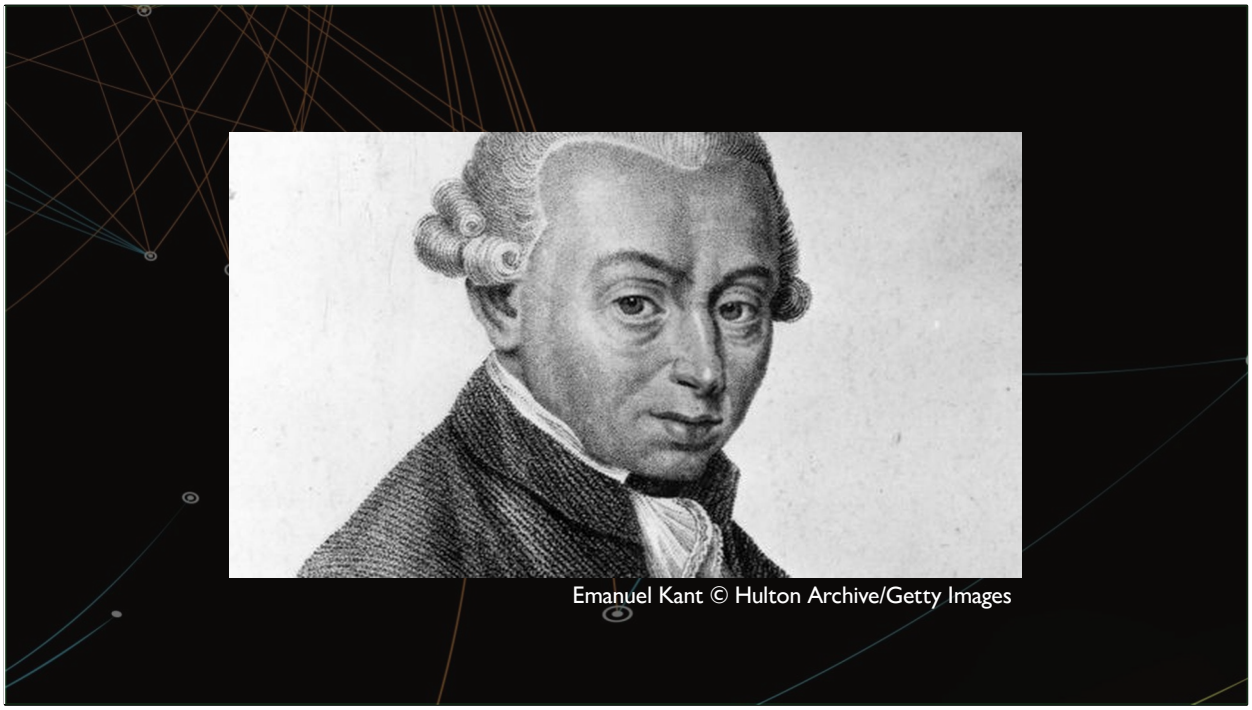

Emanuel Kant © Hulton Archive/Getty Images

Let's talk about Emmanuel Kant and the Categorical Imperative.

The Categorical Imperative is a set of rules devised by philosopher Emanuel Kant in 1785 to evaluate actions for the assessment of ethics. The rules he prescribes are **universal** and **impartial:** they does not prescribe set behaviors that you should follow, instead offering a process for testing if something is unethical.

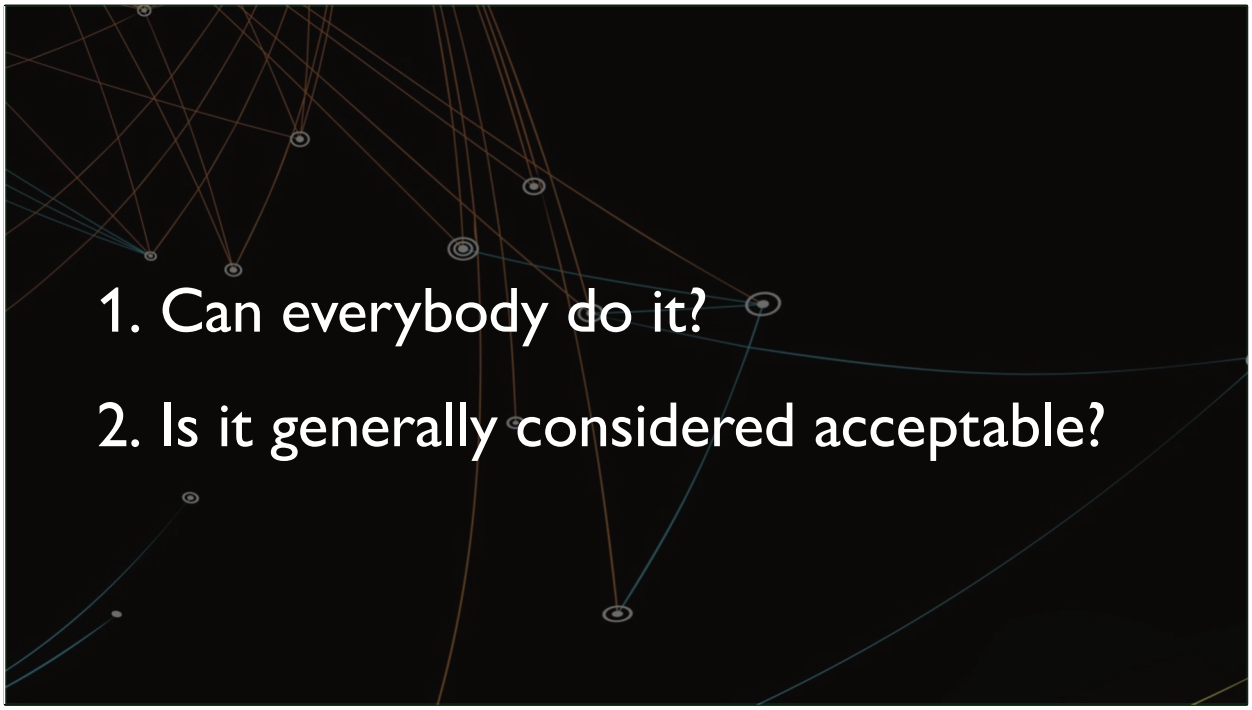

This is not a discussion on philosophy, but we can simply the rules of the the Categorical Imperative in two tests:

- Is it generally acceptable?
- Can everyone do it?

Let's look at a couple examples of applying these rules to judge if something is ethical.

Example 1: Is it ethical to lie?

Telling a lie is unethical, because it is not only generally considered unacceptable, but also because everyone can't lie since that would eliminate trust. This seems straightforward, and an easy application of the Categorical Imperative.

Example 2: Is it ethical to buy, read, and return books on Amazon?

This is something I think people would generally find to be unacceptable behavior, but even if it were acceptable, the returns are charged back to the authors, and it everyone did it authors would not be able to make a living writing.

Let's look at another example of applying the Categorical Imperative: stealing from the rich to give to the poor (Robin Hood).

Example 3: Is it ethical to steal from the rich to give to the poor?

Kant would say it is unethical because if everyone did it, there would no longer be any rich people to steal from. This is contrary to a lot of popular opinion that might say otherwise, justifying that it is right to steal from those who have excess to give to those who are in need. (I take this insight from conversations with college students studying ethics who often contest that Robin Hood behaved unethically.)

Question: Can we use the Categorical Imperative as a tool to judge our cyber security actions? Let's try by looking at an example.

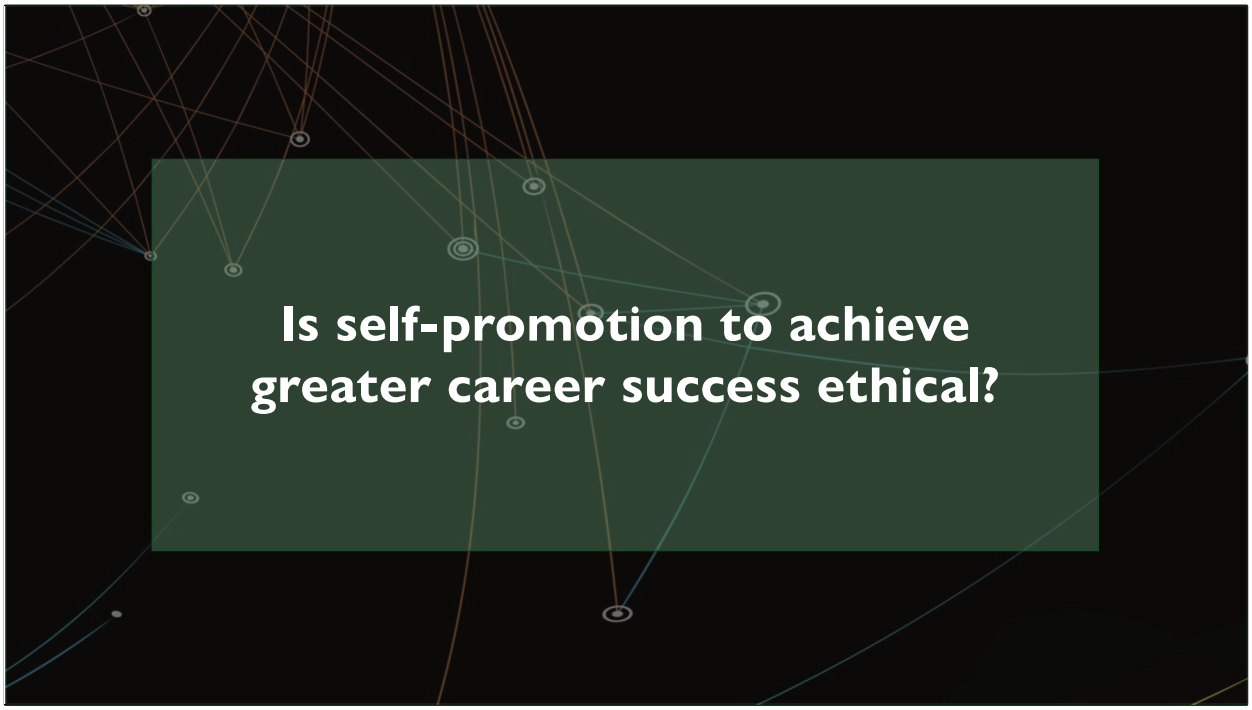

In the cyber security field, if you want to be successful, you need to be visible. This visibility can come in multiple forms:

- You can share ideas
- You can publish books, papers, and articles
- You can speak at conferences
- You can write and publish security tools
- You can talk to the press and offer expertise for widespread dissemination

Is self-promotion to achieve greater career success ethically justifiable? It's dangerous to justify something as ethical based on who the beneficiaries are for an action (the ***utilitarianism*** approach), but we can use it to evaluate immorality. If I look at my own actions, I can ask *who benefits from my actions to promote myself?* I can think of several:

- Cyber security community benefits
- End-users benefit with improved security
- I benefit
- Bad actors benefit

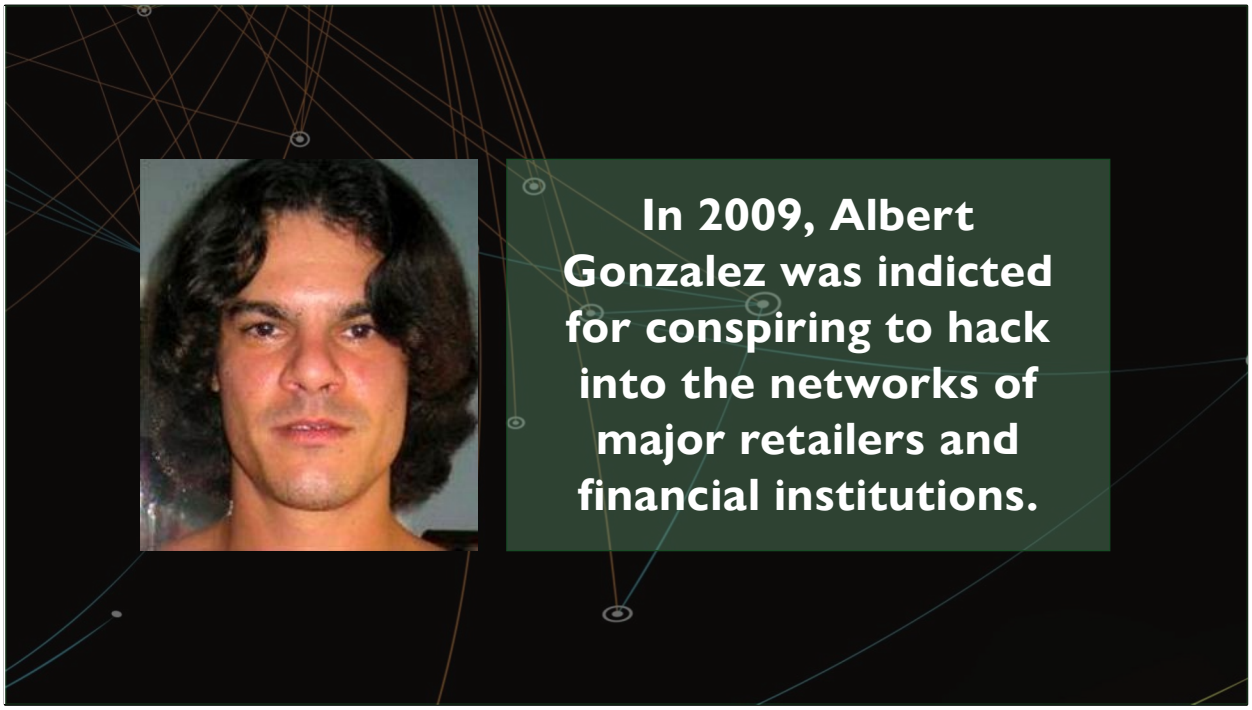

In 2009, Albert Gonzalez was indicted for conspiring to hack into the networks of major retailers and financial institutions in the US, and for the theft of hundreds of millions of credit card numbers.

In the published indictment, multiple major retailers were identified as victims:

- TJ Maxx
- Heartland Payment Systems
- 7 Eleven
- Boston Market
- DSW
- Forever 21
- OfficeMax
- Sports Authority
- Barnes & Noble
- Dave & Busters

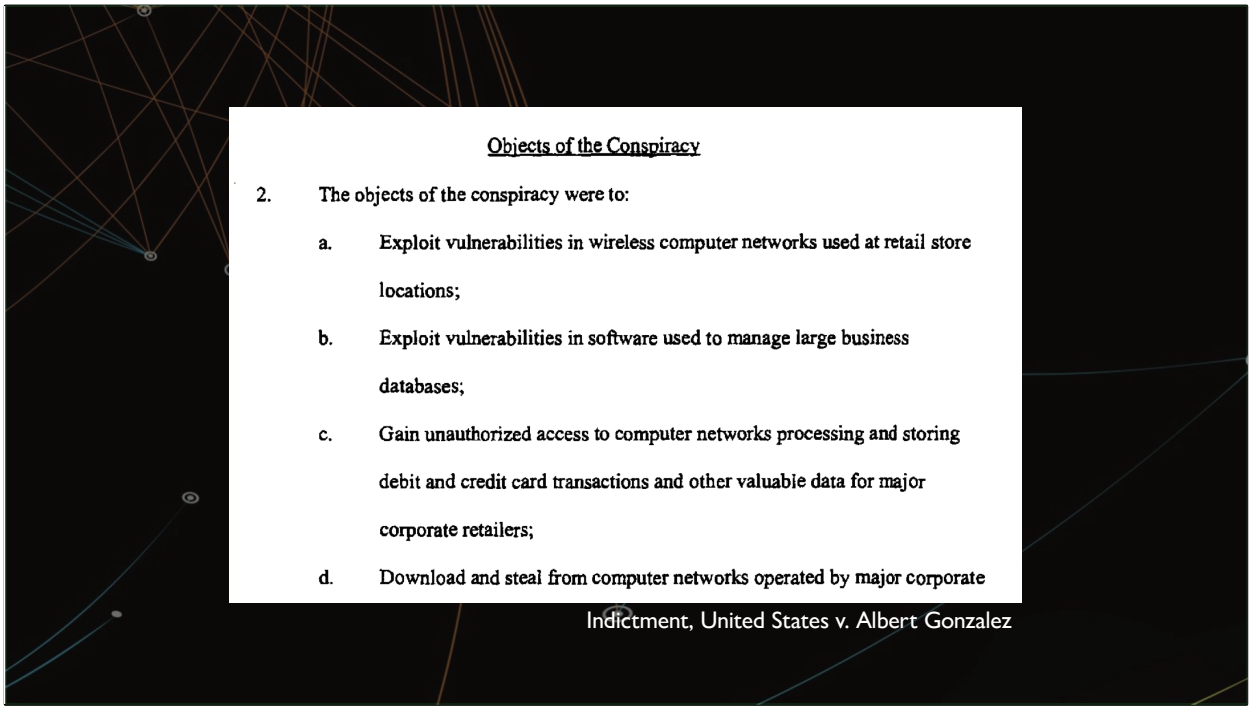

The indictment report (https://fm.cnbc.com/applications/cnbc.com/resources/editorialfiles/2012/04/19/AGS5-E40C1-SKMBT-C45110120712520-v1.pdf) indicates that Gonzales used "readily available" hacking tools to exploit deficiencies in retailer Wi-Fi systems. From the indictment:

> *The objects of the conspiracy were to:*
>
> - *Exploit vulnerabilities in wireless computer networks used at retail store locations;*
> - *Exploit vulnerabilities in software used to manage large business databases;*
> - *Gain unauthorized access to computer networks processing and storing debit and credit card transactions and other valuabie data for major corporate retailers;*
> - *Download and steal from computer networks operated by major corporate institutions*

Wi-Fi hacking tools used by Gonzalez included Asleap and Cowpatty.

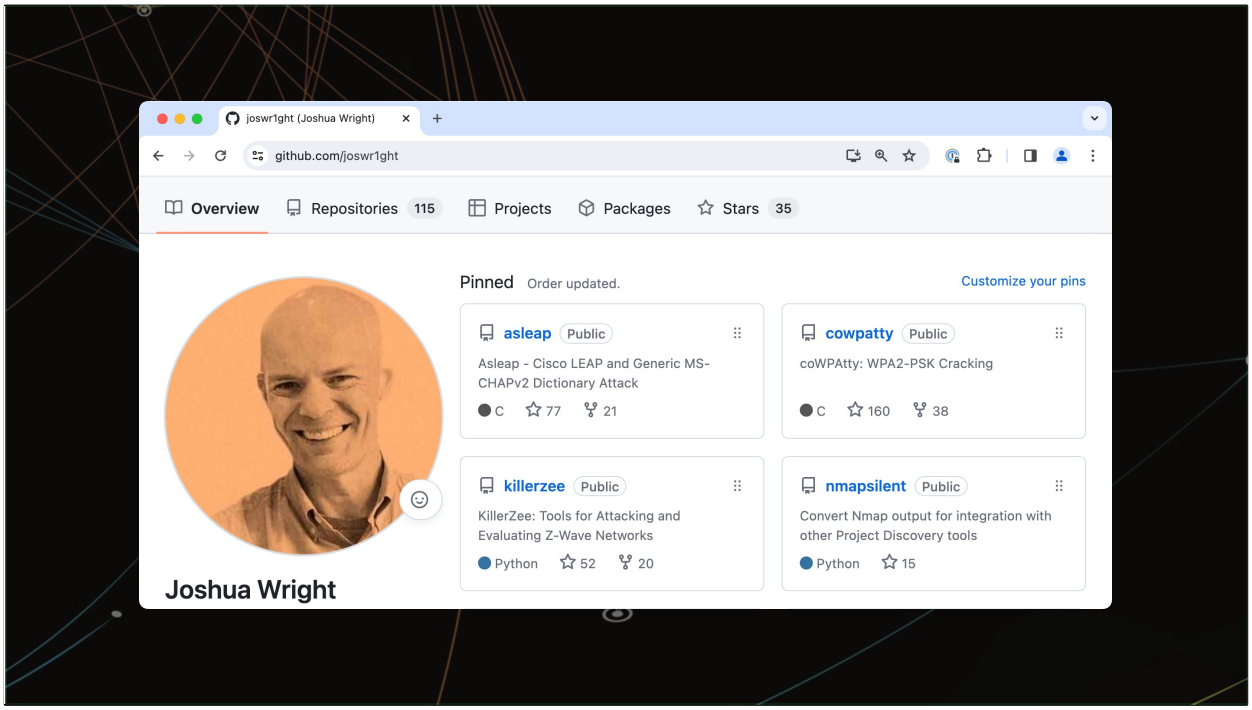

I wrote those tools.

While I knew that it was *possible* for bad actors to use my tools to commit crimes, it was easy to dismiss. Up until I learned that Gonzalez used my tools as part of his crimes, it seemed like a remote possibility, one that I could ignore to believe that I was improving security by forcing Cisco and other vendors to improve their security.

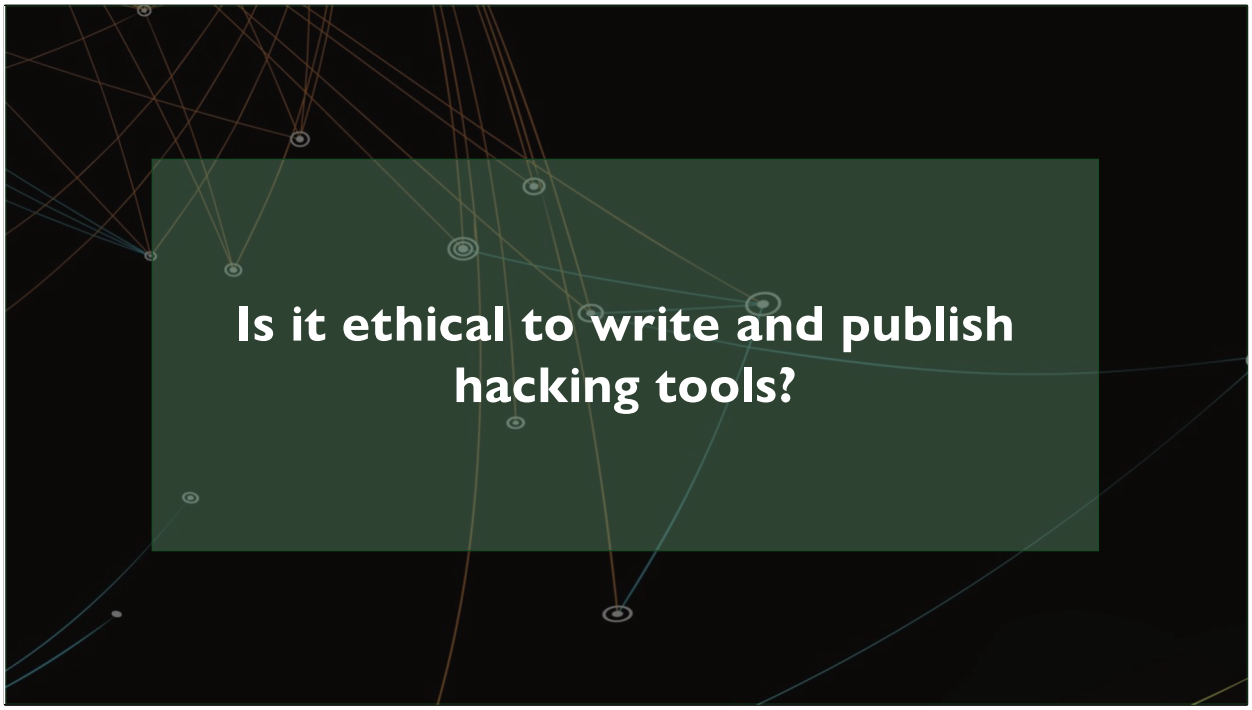

Is it ethical to write and publish hacking tools?

If we apply the rules of the Categorical Imperative, I believe it is possible that everyone could write and publish hacking tools, and I also believe that it is generally considered acceptable for people to do so within the cyber security industry. Even in 2024 however, there is a lot of discussion and criticism of people who publish Offensive Security Tools (OST).

I try not to consider the idea of *if I didn't publish the tool, someone else would have* – I feel like that is a cop-out. We need to evaluate our own actions independent of what other people may or may not do.

In hindsight, I let my motivation for self-promotion (peer accolades, speaking opportunities, credibility building) overshadow the real threat of exploitation, theft, and the impact to many organizations. I like to think that I disclosed vulnerabilities ethically (e.g., by communicating threats and waiting for vendors to make new protocols available), but I don't always feel that way each time I ask myself the question.

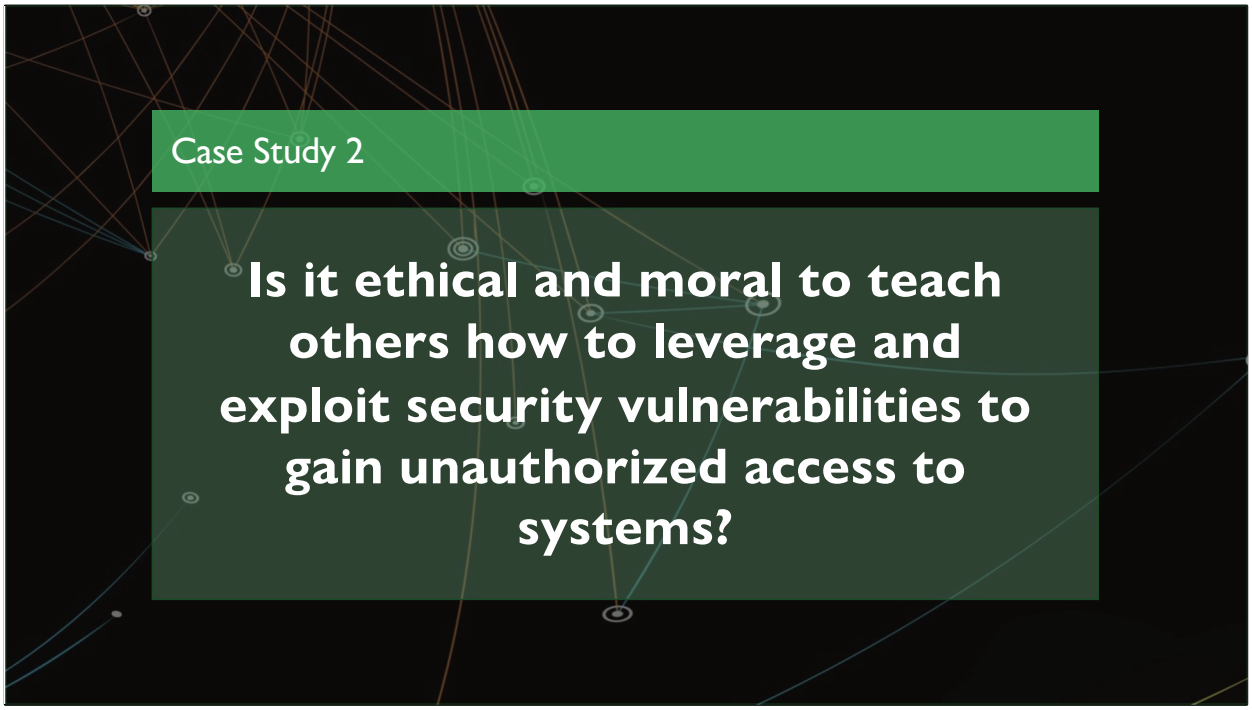

Case Study 2

For 20 years I have written courses and taught as an instructor for the SANS Institute. Most of my work is in the offensive operations curriculum, where I've written materials for advanced penetration testing, wireless hacking, and mobile device hacking. I've taught at large and small conferences, I've been a multi-year *highly rated* instructor at RSA, and I've spoken at multiple hacker conferences, sharing what I've learned openly to inform and inspire others to build on what I've learned.

Most of the time I feel as if I've made a positive difference in the world, teaching offensive security techniques. But not always.

Is it ethical and moral to teach people how to leverage and exploit security vulnerabilities to gain unauthorized access to systems?

First, let's talk about the important distinction of ethics vs. morals.

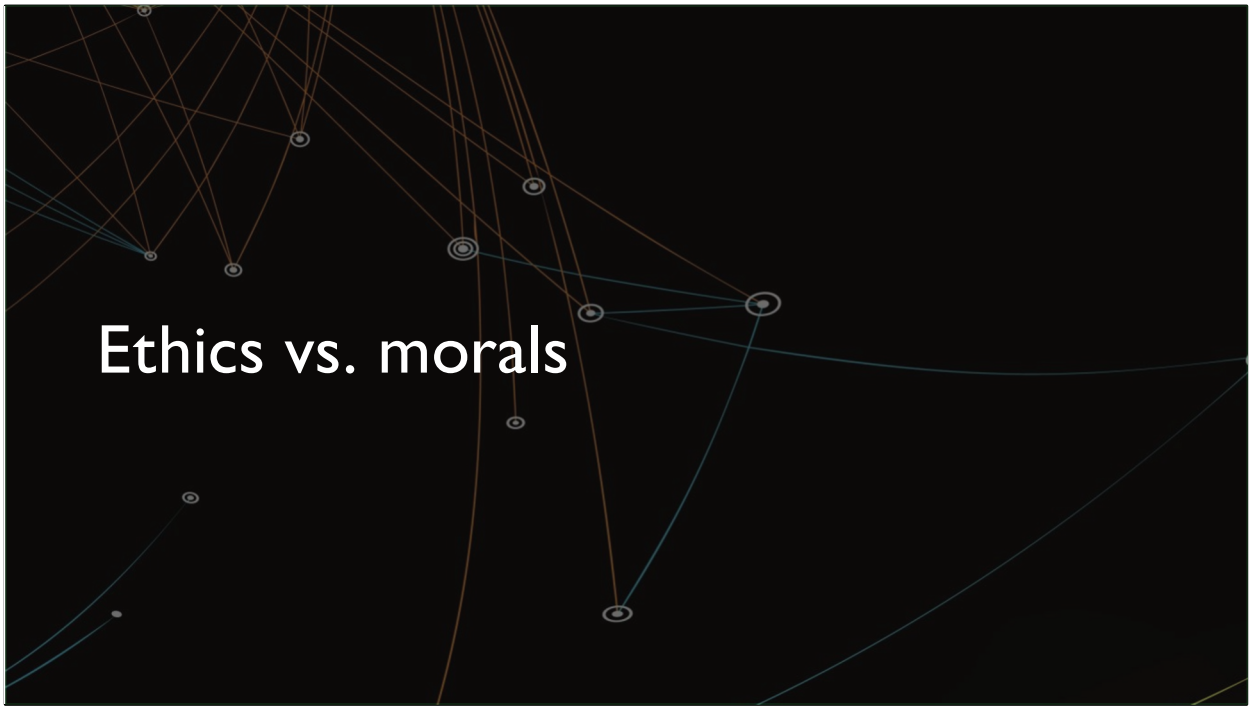

Narrowly, ethical behavior might be defined by your personal behavior, such as respecting the privacy of customer information. From a business perspective though, OFAC sanctions (https://ofac.treasury.gov/sanctions-programs-and-country-information), Know Your Customer (KYC) rules, and Anti-Money Laundering (AML) requirements are business requirements that speak to ethical behavior – they are rules and expectations for the business. Failure to comply with these expectations can be unlawful, and unethical.

However, just because something is lawful and seen as ethical by some, does not mean that it is also *moral*. Ethics are judged externally, by using the Categorical Imperative, or by reviewing common obligations such as compliance with laws or contracts, by following the common expectations for ethical behavior by people in the cyber security industry, or by reviewing the stated Rules of Engagement for a penetration test, for example.

Moral consideration can only be judged by you. It is an internal consideration, where you alone can evaluate if an action is moral.

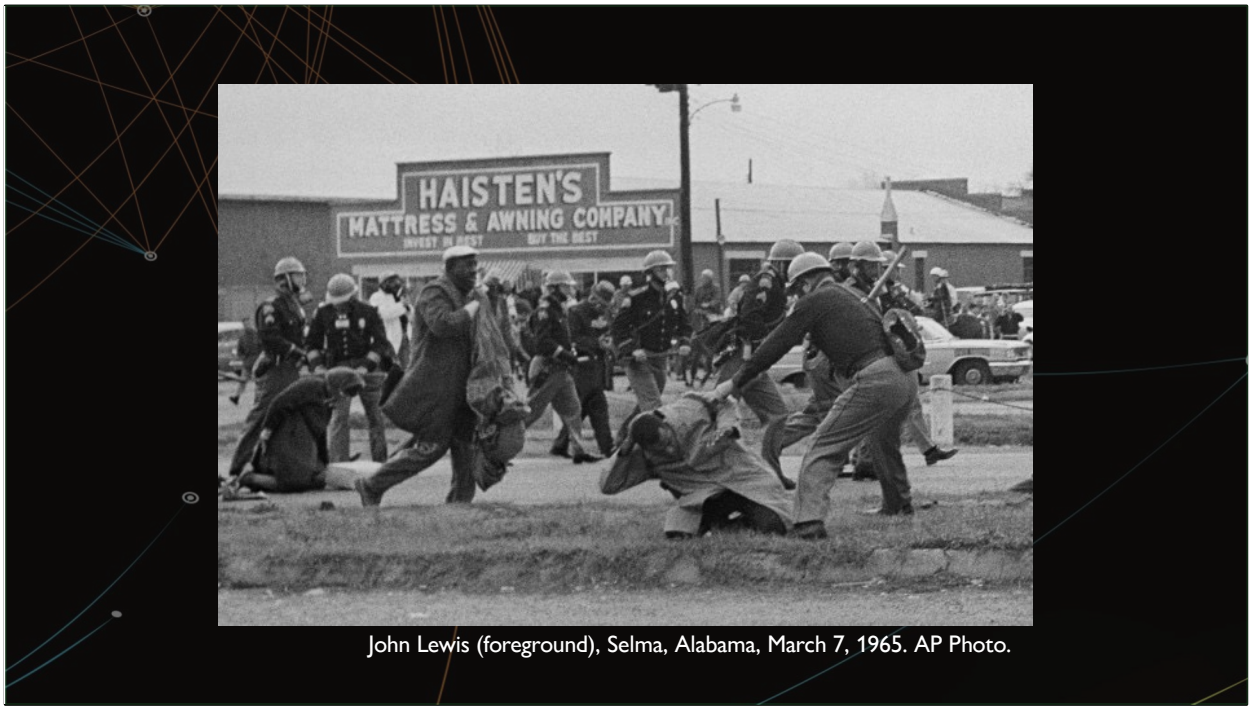

John Lewis (foreground), Selma, Alabama, March 7, 1965. AP Photo.

Several years ago, a mentor and I were discussing ethics and he asked me (rhetorically) if I thought the Bloody Sunday (https://www.politico.com/story/2018/03/07/this-day-in-politics-march-7-1965-437394) attacks by police on civil rights protestors marching from Selma to Montgomery were ethical. I said *yes*.

Since the police were authorized by county officials to break up the marchers in violation of Judge James Hare's injunction making it illegal for more than two people to gather under the sponsorship of civil rights organizations, the police were acting in compliance with the law. I told him I thought the police actions were ethical in a strict interpretation but were also *morally reprehensible*.

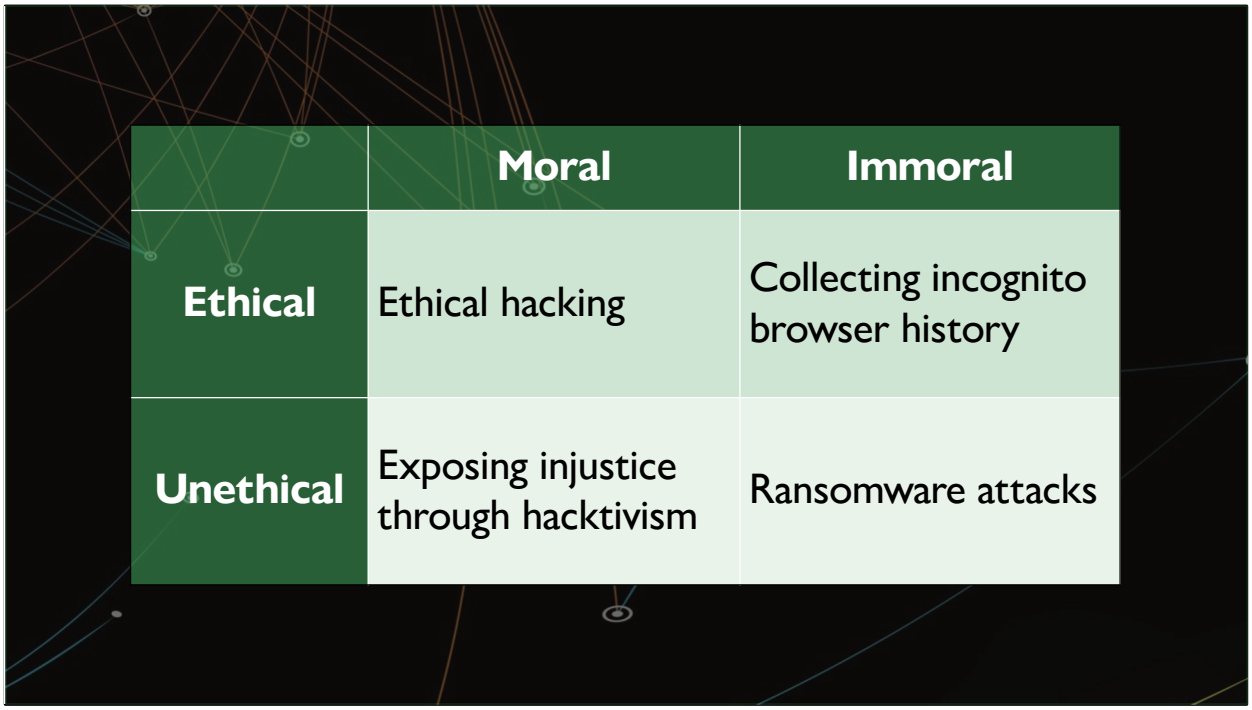

When we consider our actions, it's useful to separate whether something is *ethical* and if it is *moral*. I've come up with some examples where something is ethical but immoral, and vice-versa, but this reflection on morality is solely my judgement. You may decide that something I believe to be immoral is moral, and vice-versa.

Let's put this understanding into practice.

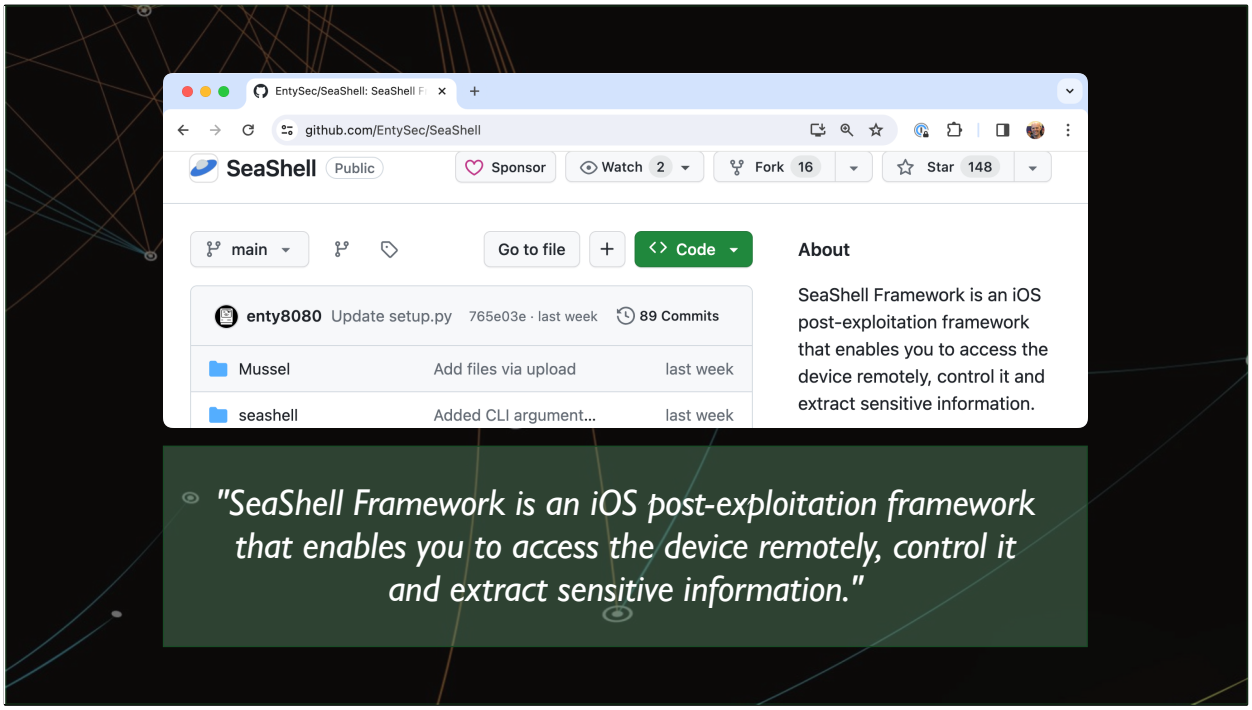

The single most rewarding part of my career is helping others to pull back the blanket of FUD to peek into the often-fragile reality of cyber security. When I'm working with students, seeing them *connect the dots* on risk through understanding, and knowing that their perspective is forever changed makes me feel like I am making the world a better place, a little at a time. I love that feeling, and it motivates me to improve my skill as an author, as an instructor, and as an offensive cyber security practitioner.

That is, until I don't.

It isn't controversial that tools can be used for good or for evil. The difference is in the wielder – how will you apply what you've learned? Are you going to evaluate your actions to ensure they are ethical?

When I wrote the SANS SEC575 course, Mobile Device Security and Ethical Hacking, I devoted a day of content to understanding Remote Access Trojans (RATs) like SeaShell (https://github.com/EntySec/SeaShell), designed to take over a victim's mobile device. Students would examine several RAT platforms, both open-source and commercial tools, and students would deploy one or more RATs on target mobile devices to exfiltrate SMS messages, email, contact information, GPS location data, photos, WhatsApp messages, and more. We would review opportunities for RAT

deployment, including using Android and iOS-specific features to sideload an application, and mechanisms where an attacker can preserve persistence on a mobile device even after a factory reset.

After class one day, two students approached me and asked if they could discuss some specifics of RAT deployment. Specifically, they asked:

What would you recommend for undetectable RAT deployment on dissident journalist mobile devices?

At first, I thought I misunderstood the question and asked if they could elaborate. One of the students continued to explain that they signed up for the class to gain insight into techniques to collect data from journalist mobile devices. Specifically, he wanted to know about collecting location history with timestamp information, contact book information, messages, and email from devices his employer identified as *persons of interest*.

Looking at the attendee badges I recognized their employer's name as a government defense organization. Awkwardly, I tried to avoid answering the question. Later I wondered why it felt awkward and uncomfortable to talk about a specific scenario where *ethically* the laws in the country might allow for this type of civilian spying, even if it's something I talk about from a non-specific example in the course material on a regular basis.

I think the awkwardness was because they presented me with an example where activities might be ethical, but I felt like the activity (violating the privacy of journalists identified as dissidents through mobile device hacking) is immoral.

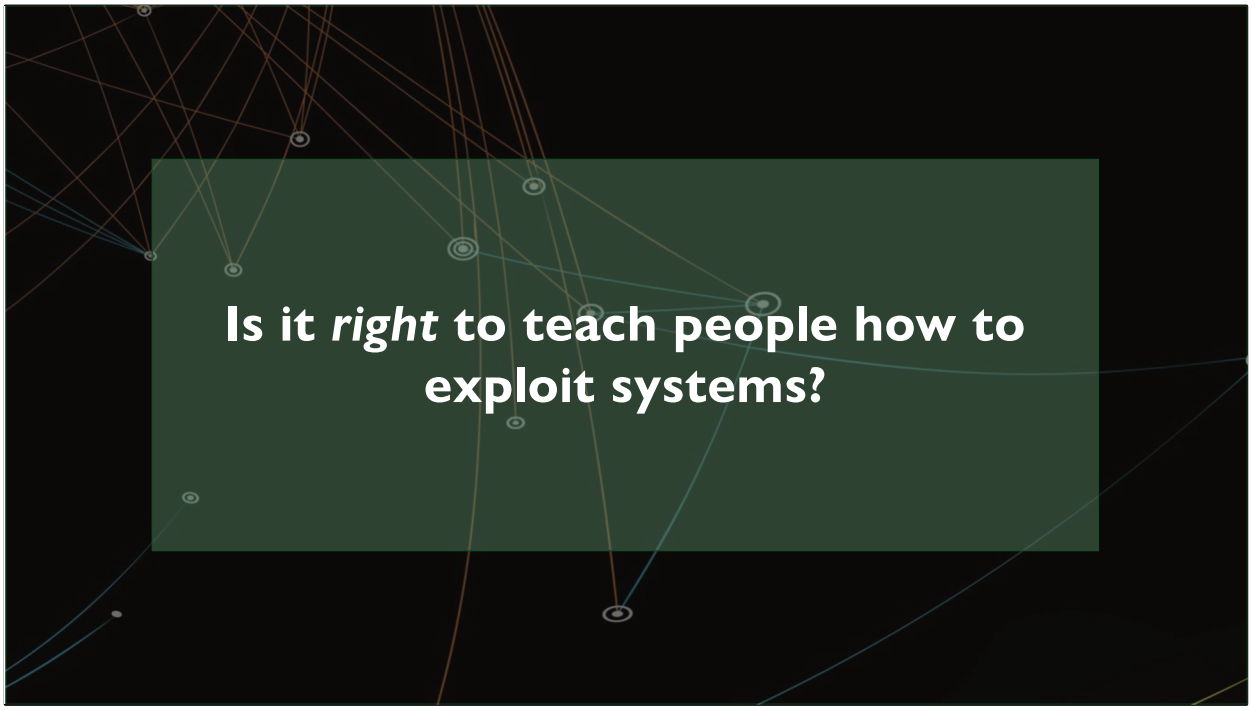

Is it right to teach people how to exploit systems, knowing that they can and will apply it for goals that I think are immoral? Am I responsible for the actions of someone else who chooses to apply what I teach in an unethical or an immoral manner?

I think teaching offensive security topics is broadly considered ethical. If we apply the Categorical Imperative, we can deuce that everyone can teach, and it is broadly considered acceptable. While *ethical*, evaluating if it is also moral is a more complicated question.

Most of the time I feel as though what I teach does more good than bad, but that's difficult to assess in an empirical manner. The most challenging moral questions are the ones where I don't know the answer each time I ask the question.

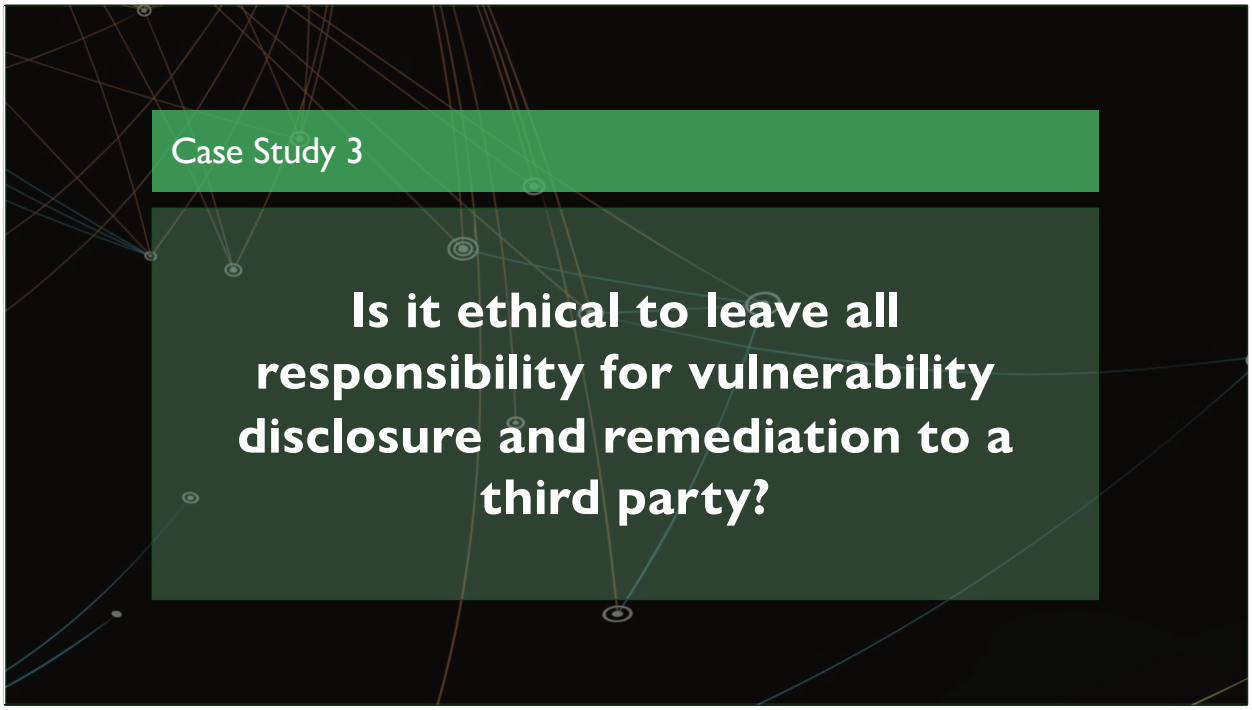

Case Study 3

Is it ethical to leave all responsibility for vulnerability disclosure and remediation to a third party?

In my career I've worked with many organizations to test the security of their networks, applications, infrastructure, products, and systems. As a professional penetration tester, my job is to identify and exploit the vulnerabilities to not only point out but to *demonstrate* risk.

For this job I am hired and paid. Sometimes it seems like a dream, reenacted in movies or TV shows. I identify the problems, I write and deliver a report, and the customer uses that report to improve their security. Everyone benefits!

But not always.

Is it ethical to assess and identify security vulnerabilities in a customer's product, and leave all responsibility for vulnerability disclosure and remediation to the customer?

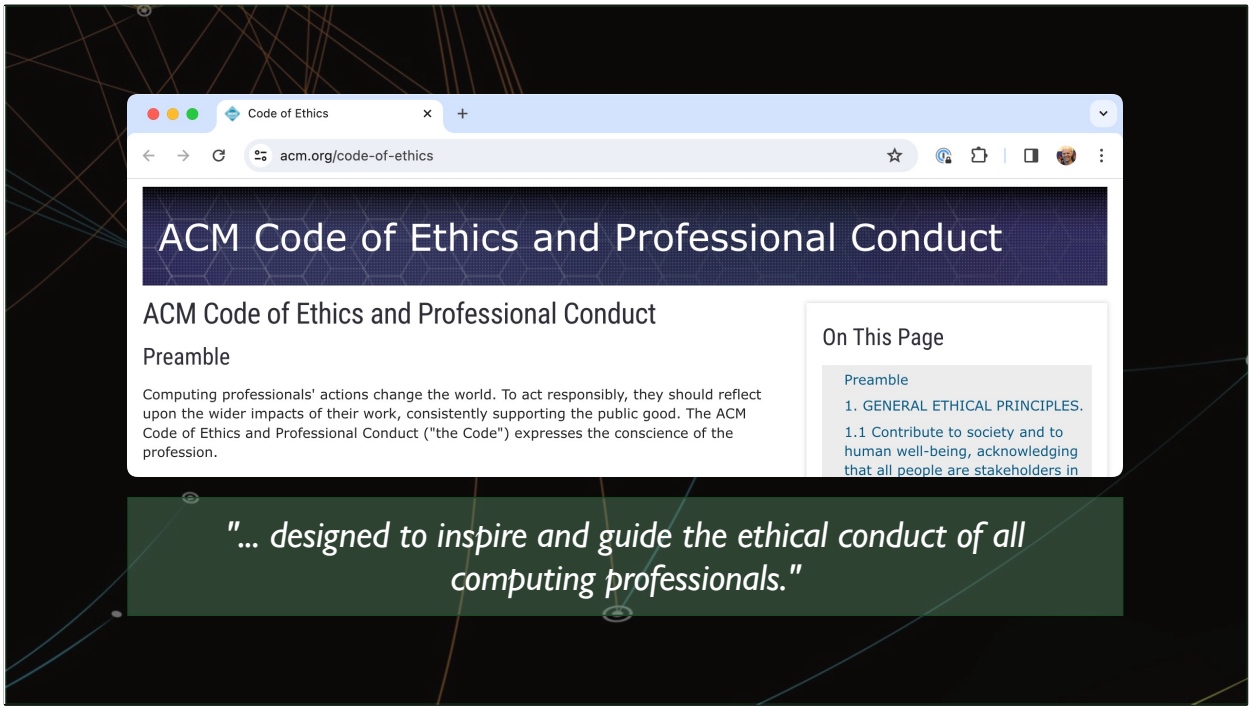

While the Categorical Imperative is universally valuable, it can be helpful to have specific guidelines that define ethical behavior. The Association of Computing and Machinery publishes a 28-page guide, the Code of Ethics and Professional Conduct (*the Code*, https://www.acm.org/code-of-ethics), "designed to inspire and guide the ethical conduct of all computing professionals."

The ACM Code defines general ethical principles, professional responsibilities, leadership principles, and compliance expectations. None of the Code's rules are particularly controversial:

- Respect privacy
- Know and respect existing rules pertaining to professional work
- Ensure that the public good is the central concern during all professional computing work

The ACM Code is a valuable resource, is well thought out, and can be used by cyber security professionals to guide decision making.

But what do we do when the rules conflict with each other?

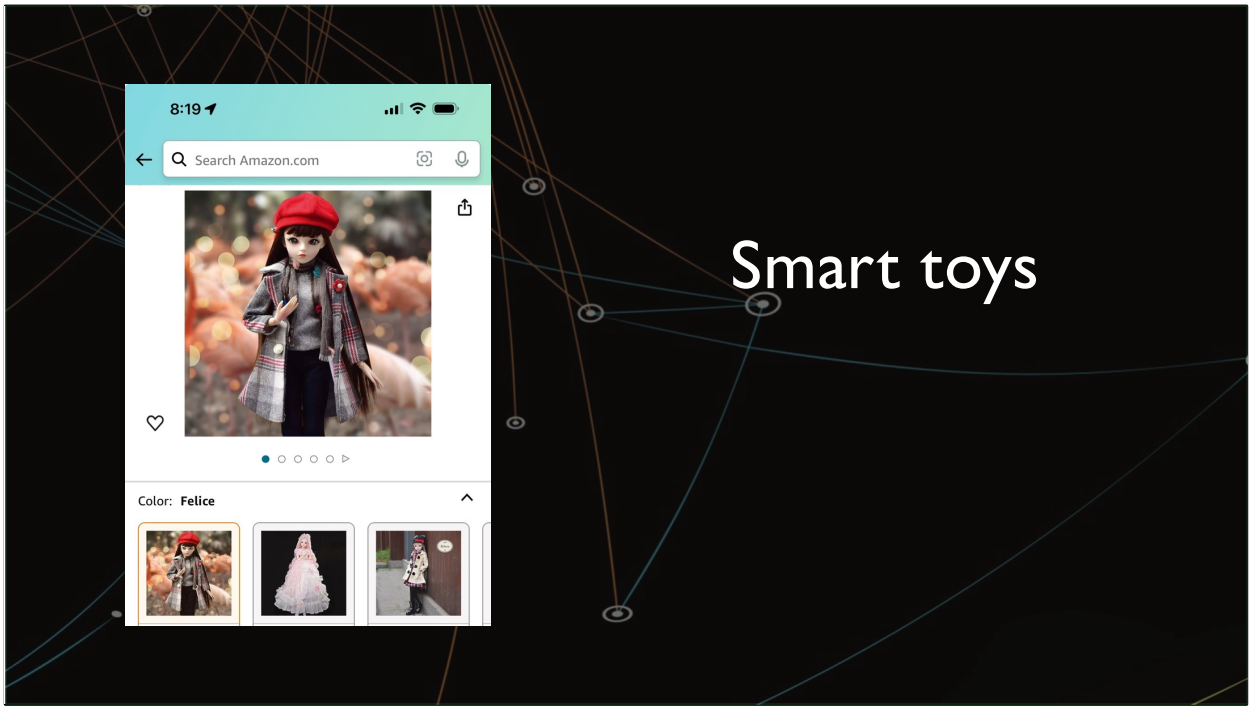

Some of my favorite work as a penetration tester is to evaluate the security of *products*. This could be a unique software package or a service that I need to evaluate, identifying vulnerabilities so my customer can fix the flaws. Sometimes, the product is a tangible item I can hold, such as a wearable device, a smart electricity or water meter, or even a smart toy.

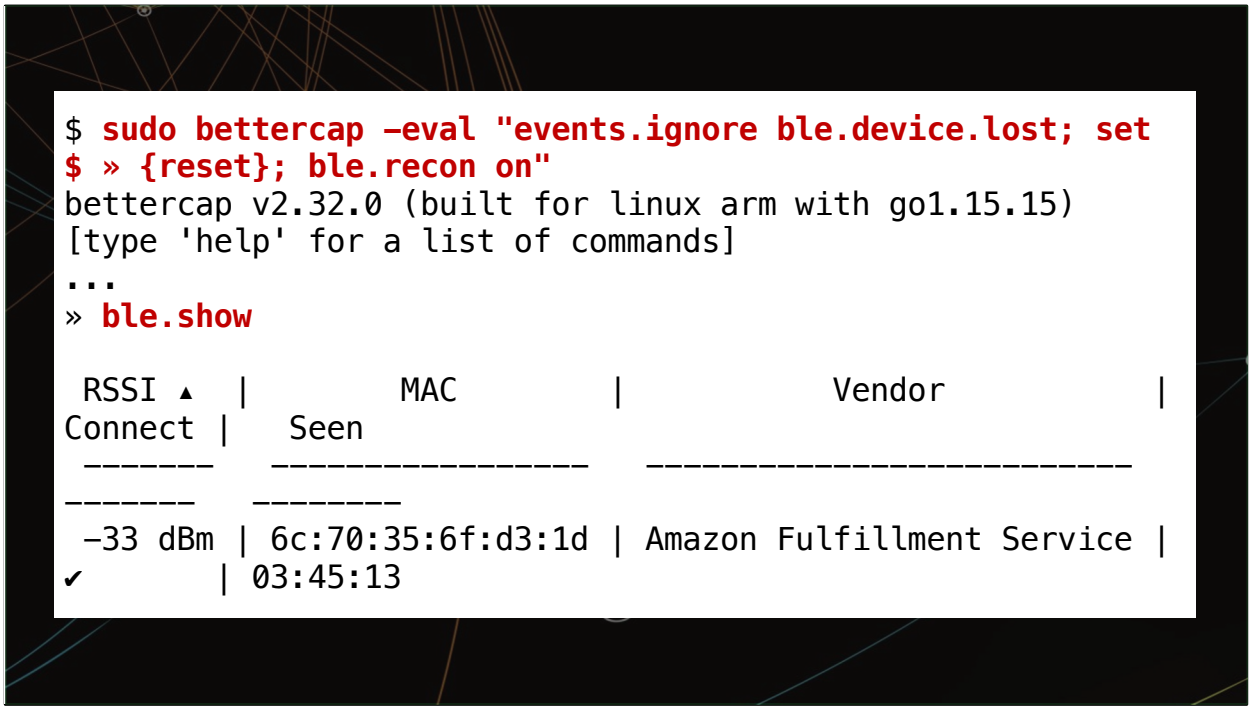

Smart toys often accommodate wireless connectivity over Bluetooth, Wi-Fi, and other protocols. In one engagement, I used Bettercap (https://www.bettercap.org/) to scan and identify the Bluetooth Device Address information for the target smart toy.

Normally, Bluetooth devices will provide connectivity to a control interface, such as a mobile app, that allows the child or parent to interact with the toy in creative ways. For this product, the parent could *speak* through the toy, using the doll's voice to render text-to-speech messages to the child.

For this smart toy, while the cloud to app security was strong and required authentication, the app to toy Bluetooth connection was not protected. After reverse-engineering the custom Bluetooth protocol between the toy and the app, I discovered that I could send custom messages to the toy, rendering any text as speech.

```
pi@raspberrypi:~ $ sudo scapy -H
>>> t = 'Can you brush my hair?'
>>> d = hex_bytes('02010610ffff') + struct.pack('>H',len(t)) +
t.encode('utf8')
>>> p = HCI_Hdr() / HCI_Command_Hdr() /
HCI_Cmd_LE_Set_Advertising_Data(data=d)
>>> bt = BluetoothHCISocket(0)
>>> bt.sr(p)
Begin emission:
Finished sending 1 packets.
```

While the application interface protected access to the audible speech interface on the toy, someone with close physical proximity to the toy could prompt it to say custom phrases. I envisioned this as a safety risk to the child, since a predator could potentially direct the child to take dangerous actions through the voice of the smart toy. I reported the vulnerability to my customer as a significant risk, demonstrating the

problem.

Later, I learned that the customer did not resolve the vulnerability, shipping the product in time for Christmas buying season as tested.

If we look at the ACM Code, we see guidelines that we are supposed to apply to direct our behavior. Does the risk acceptance for a child's toy warrant violating these rules? Do I have an additional obligation beyond my commitment to respect my customer's privacy and respecting the rules of professional work (honoring NDA)? If I feel that public good is not the customer's priority, do I have an ethical obligation to report my concerns to a higher authority within the company or even a law enforcement agency?

Rules are useful in governing ethical behavior, but they can be reductive too. Ethical assessment can be more complicated than what we can easily resolve in a 28-page PDF.

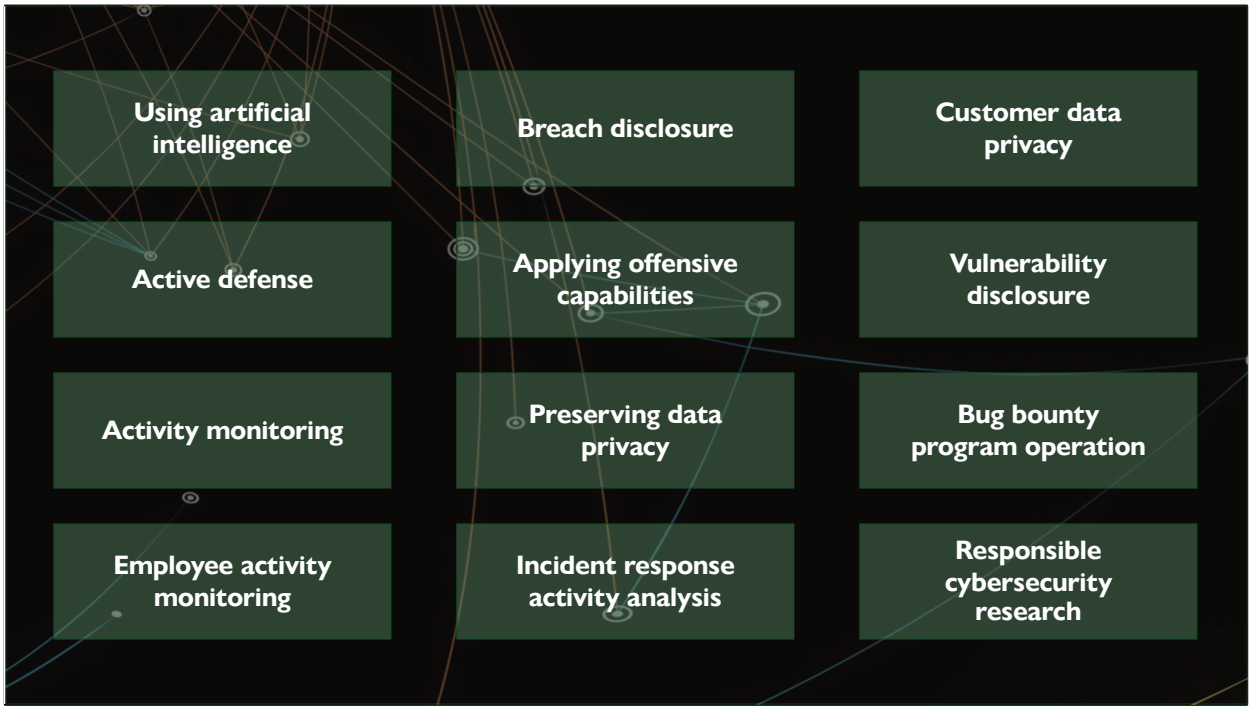

The ethical challenges we face in cyber security are broad. As professionals, we should keep asking ourselves if we are acting in a manner that is ethical and moral. This can be, and should be difficult, requiring us to collaborate with others to review and assess our actions.

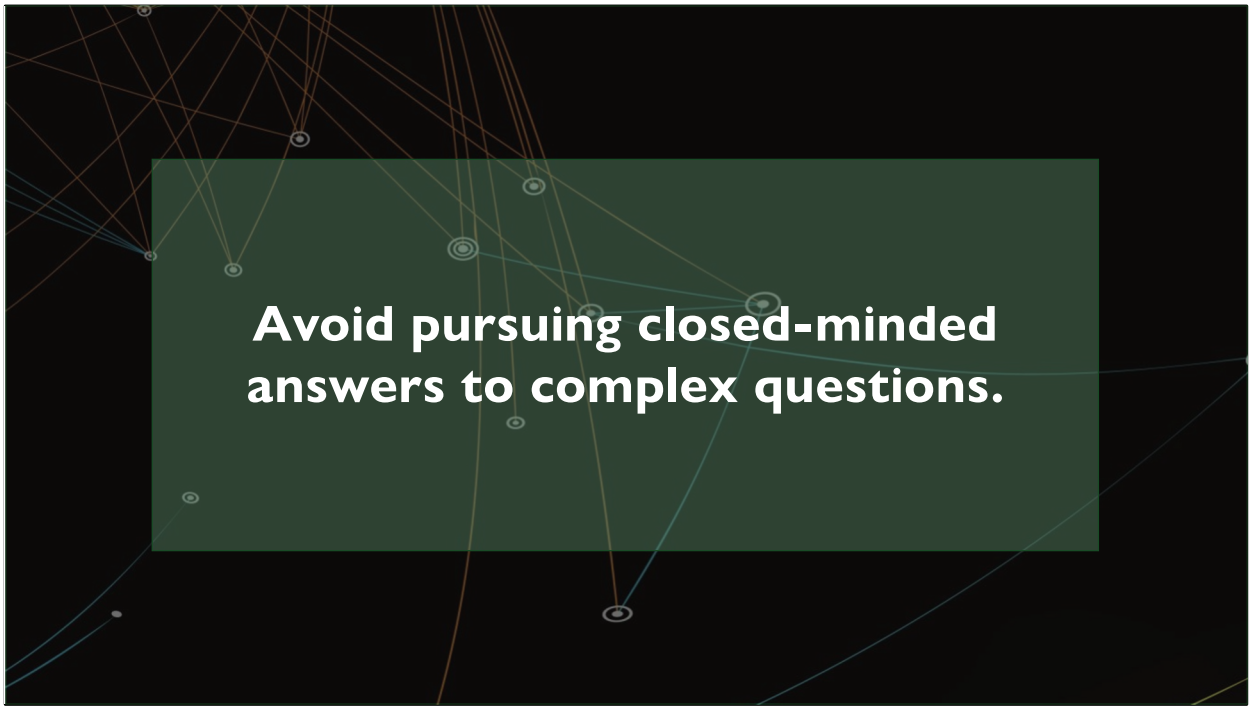

Each of the case studies we look at in this presentation ask closed questions to pursue an ethical/not ethical answer, but our work in cyber security is often much more complicated than that.

Security analyst Chris Sanders wrote an excellent article on the Offensive Security Tools (OST) debate (https://chrissanders.org/2020/07/research-ost-release/) where he offers suggestions to help us change our perspective. Instead of asking "Is it right to release OSTs?" he asks "**To what degree does the release of offensive security tools make the world inherently more secure or insecure?**" (paraphrased)

When we ask questions that have a binary yes/no, ethical/unethical answer, we remove all nuance and create oppositional forces where people on either side of the question want to be *right*. Instead, we need to consider redefining out questions in ways that are open-ended such that they can be debated and considered without creating opposing sides. In this way, we encourage open discourse, and we create a broader opportunity to see other people's opinions and insights to help us consider ethics and moral responsibilities.

It is tremendously valuable to have a relationship with a mentor that can offer guidance, insight, and even questions you may not have considered for ethical challenges. Ideally, this mentor should be outside of your immediate social circles too, since we often form our social circles with people who share the same ideals and opinions as us.

Something that has helped me grow is to review big decisions I have made in the past where I was challenged by ethics. I always ask myself if I would make a different decision today, using what I knew at the time, but also what I know now and perhaps should have considered.

While there are challenges with ethical behavior guides (like we saw in the ACM Code of Ethics and Professional Conduct) that can be challenging, it is still useful to adopt and apply a guide specifically written for the cyber security industry. Some options include:

- ACM Code of Ethics and Professional Conduct (https://www.acm.org/code-of-ethics)
- UK Cyber Security Council Ethical Declaration (https://www.ukcybersecuritycouncil.org.uk/ethics/ethical-declaration/)
- Montreat College Cybersecurity Oath (https://www.montreat.edu/cyber-oath/)

These guides should be applied as an set of infallible rules, but rather to help inform you as you evaluate your own actions, taking into consideration the collective wisdom of those who assembled them.

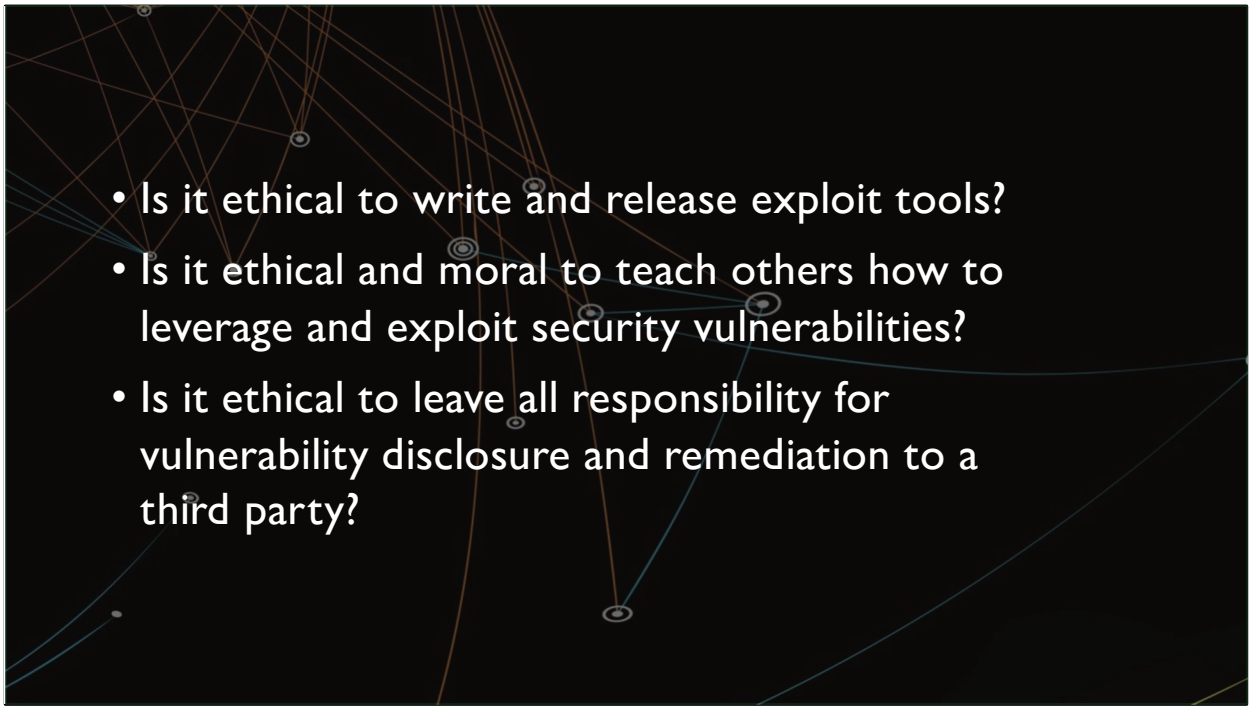

I don't know that I've always made the right decisions in my career from an ethics perspective. I try to make just decisions by applying standardized and specific frameworks to guide ethical behavior, but sometimes I look back and think I made decisions in haste.

Ethical decision making is challenging, as it should be. If we aren't challenged by ethics, or if we believe that we've never made a mistake regarding ethical decision making, then we're probably not looking closely enough at our actions. The process of ethical consideration is always warranted, to help us guide our actions.

Even when I feel like I made poor decisions, or behaved in a way that is unethical or immoral, I take the outcome as an opportunity to do better. And I take assurance in knowing that since I keep asking those questions, I am demonstrating ethical behavior in decision making, while I work to guide my future decisions.


Thank you.
Joshua Wright
Faculty Fellow, SANS Institute
josh@wr1ght.net
@joswr1ght
Many thanks to those mentors and colleagues who have helped me shape and examine my professional ethics and morals:
Al Benoit
Anwar Abeyie
Chris Sanders
Dr. Rev. H. Daehler Hayes
Ed Skoudis
Jared Folkins
Jennifer Wright
John Strand
Judy Novak
Kevin Finisterre
Merwyn Andrade
Mick Douglas
Mike Poor
Rev. Joy Utter
Rob Lee
Stephen Northcutt