Chapter 1
Why Spy on Bad Guys?

I am sure you are wondering why the *hell* a commercial cyber espionage operation is necessary. I assure you that it is not because doing so is fun. Don't get me wrong—it is—but that isn't why we do it. We do it because the industry needs us; our clients need to know what the bad guys are doing so they can plan a defense. It was a rough evolution to get here, though.

I once consulted for a chief information security officer (CISO) in the retail space. He had informed me of his investment in a software-as-a-service (SaaS) cyber-intelligence vendor. I inquired if I could ask a few questions about his program. Even though it was clear the word *program* confused him, he agreed.

"Okay, so you have some very specific use cases for intelligence—is that why you licensed this platform?"

"Um, not specific, more general."

"Okay, but you know what specific data your company needs to turn into intelligence and for whom inside the organization?"

"I think so . . ."

"Okay, so you are hiring a full-time intelligence professional to sit in front of this thing?"

"No."

"Oh, so you are taking one of your existing staff, training them on intelligence process and production, and putting them on this full time?"

"No."

"So you are taking the person that reads your EDR [endpoint detection and response] events all day and letting them query this platform when they have time?"

"Kurtis, you are making this sound like it's not a good idea . . ."

"That's the point. But let's say this all works well—it won't—but hypothetically, let's say it does. And your EDR person runs a query one day and finds a bad actor on the dark web talking about a backdoor into your online system where they can change the prices. They are selling the price change as a service to anyone who wants to purchase products from you—that is, someone pays them a small amount in Bitcoin, and then they change the price of a very expensive product to $1."

"Yes! Yes! This is exactly the kind of intelligence I am looking for!"

"That isn't intelligence, but I see. Okay, next question. What do you do next?"

"Um . . ."

"You see, there are many bad actors on the dark web. Most of them are full of shit. How do you validate this before you start ringing alarm bells? My guess is you have *a lot* of APIs. Which one is it? What is the vulnerability that was exploited? Is this bad actor credible? Can you see if they have already received payments? Can you talk to them?"

"I don't know . . ."

"Then what good is this platform? If it just sends you data or information that you cannot validate and cannot act on, that sounds like a terrible investment."

This is how it goes, and so it is with intelligence platforms at the enterprise level. So many of the suppliers and vendors are selling a tool that produces data for organizations that need intelligence. Further, most organizations cannot act on the data, information, or intelligence they receive. This is why it is critical not to look at cyber intelligence as a product but rather as a program. It is a supporting

business operation and requires process and talent like any other business function.

This chapter aims to baseline what cyberthreat intelligence is and *why* we do it. We will delve into a bit of the *how*, but the details of *how* are covered in Chapter 3, "Cyber Espionage 101." Chapter 4, "Whodat?" will cover *where*.

The evolution of the cybersecurity product and service market has been largely unbalanced, its offerings primarily reactive in nature and typically from a defensive position. Over the decades, the cyber adversaries' sophistication, motivation, and collaborative nature have increased considerably, and the vendors have responded with more defense tools, technologies, and service offerings. The defense-in-depth cyber stack has begun to overwhelm even the largest and most forward-thinking organizations. CISOs are inundated with new gizmos and coerced into buying the newest thing to defend against the threat.

What are the real threats? Who is targeting your organization and why? What are their motives and capabilities? What tools do they have? What data about you do they have? How successful have they been attacking similar organizations in the past? Are they script kiddies or nation-state cyber armies?

The U.S. federal government has the largest military defense budget of any country globally by a significant margin, but it isn't limitless. How does the government decide how to allocate that budget? The U.S. federal government uses intelligence to inform them of their adversaries' motives and capabilities. The Central Intelligence Agency (CIA) is the tip of the spear for defining the nation's defensive (read: "offensive") strategy and budget allocation. Instead of guessing and trying to defend ourselves against virtually anything, which would be a poor strategy for our tax dollars, the CIA gives our military the information necessary to focus their capabilities. The CIA uses myriad sources to gather this information. These include signals intelligence (SIGINT), open source intelligence (OSINT), and the one we watch movies about, human intelligence (HUMINT). The CIA correlates this information in Langley, Virginia; validates; adds context; assesses the impact; and

produces intelligence. Of all the sources, HUMINT is one of the most impactful, as it is the most tangible and most easily validated source. After all, our nation-state adversaries are also human.

Likewise, CISOs have limited resources. For most of the history of the cybersecurity market, they have been building digital walls and response systems to defend against unknown attacks and ghostly attackers. Gathering information about the threat to a specific organization and its would-be opponents helps CISOs focus their spending.

To do this effectively, one has to take the fight to the bad actors. Not unlike with a CIA asset or field operative, this invaluable information has to be gathered from the trenches, the chatrooms, the dark web forums, and the private chats of those who would do us harm. It is, in fact, cyber espionage.

As a result, an ecosystem and extensive market providing such services has burgeoned. Companies have been integrating and maturing a practice to keep watch on adversaries for over a decade. Many of these companies fall under the cyberthreat intelligence (CTI) or threat intelligence (TI) umbrella. Unfortunately, there is no litmus test on how to deliver these services well, nor has there been a quantitative method for measuring success. The net result is that many of these programs serve as a box for a CISO's security program checklist, and business outcomes are elusive. One of the key reasons this continues to be a challenge is the focus on providing raw data rather than real intelligence.

Some of the earliest versions of TI were simply feeds of known bad digital indicators. There were seemingly independent operators, like Spamhaus and the SANS Internet Storm Center, who supplied lists of known bad IP addresses and domains. The security vendors jumped in around the same time in the late 1990s, often crowdsourcing their threat data from their software and appliances on their customer sites. If Firewall A at Company A detected an attack or malware from IP X.X.X.X, they would tell companies B–Z to automatically block any traffic from X.X.X.X. This was immensely valuable, and versions of this kind of threat intel feed exist today. In fact, most EDR vendors like CrowdStrike and SentinelOne share

their threat telemetry data across customer installations in a similar manner.

The next iteration in threat intelligence is thought by many to have been pioneered by iDefense, a company founded by James Adams and later run by John Watters; they are credited with pioneering the next phase of the cyberthreat intelligence industry. iDefense began providing threat actor profiles, advanced persistent threat (APT) monitoring, and data on zero-day development and activity.

John Watters' next company, iSIGHT Partners, took the iDefense playbook and advanced it iteratively. By employing hundreds of analysts across the globe, iSIGHT infiltrated and interacted with the threat actor community. They delivered detailed "ThreatScape" reports with a particular vertical or cyber domain focus to their high-paying customers. FireEye acquired iSIGHT in 2016, and remnants of the iSIGHT intelligence practice remain inside what is now Mandiant, a Google Cloud company.

By 2013, even iSIGHT's army of analysts couldn't scale to deliver what customers were beginning to ask for. Customers wanted tailored cyber intelligence for their business and their brand. What's more, they didn't want it in report form; they wanted it in as close to real time as possible.

In 2014, when two partners and I launched the company that would become GroupSense, we knew the benefits of what we were doing. The original use case was simple and narrow, with just three pillars:

1. Find the customer-specific data being pilfered by the bad guys on the Internet before they monetize it, and let the customers know about it.
2. Find bad guys talking about specific customers and tell the customers about it.
3. Find infrastructure that bad guys set up to attack specific customers—and tell the customers about it.

We worked long and hard to develop avatars or personas trusted by the underground Internet community. We were successfully invited to participate in the forums and channels where the illicit

activity occurred. We had a platform that was capturing those conversations via our personas and analyzing them at scale for customers. Jackpot, right?

Not exactly. We were lucky to have landed a couple of the largest brands in the world as early-adopting customers. Our platform was successfully identifying digital artifacts that, in the wrong hands, could cause material damage to their businesses. The software was successfully alerting them to those digital artifacts. "ALERT! Bad Guy on Bad Channel is offering to sell this Intellectual Property of yours!" The customer was disturbed but pleased with the platform's effectiveness. Yet, only weeks into their use of our product, they had begun calling us. "Hey, we firmly believe the data you sent us about the bad guy is . . . well, bad. We just don't know what to do next." This is a challenge companies all over the world struggle with when it comes to cyber intelligence and dark web monitoring programs. How does a cyber organization take action on or operationalize that data or information?

We knew what the answer was. Someone had to talk to the bad guy to determine if the intellectual property was indeed stolen, if it was real, and if it was truly a threat. The customer couldn't do that. Not only did they not have a tenured, trusted avatar on the dark web, but this advertisement was made by a Russian actor in a Russian forum, where they speak—well, Russian. To make things worse, even if they had a Russian-speaking staffer with access to a tenured sock puppet in the relevant forum, their legal department would not allow them to engage.

We had to ask ourselves if we spent all that time creating a poor product, delivering a solution that kept people up at night, and sending alerts that could not be verified or acted on.

Driven by a sense of duty and purpose, we did what we had to do. We talked to the bad actor, got sample data, verified where the leak was, weighed the potential risks of not acting, and sent a custom alert to the client explaining the situation, the context of the forum, intel on the actor and their credibility, screenshots of the sample data, and a recommended course of action based on that information. We sent them actionable intelligence in a report we later renamed an "advisory."

Much has been written about the difference between data, information, and intelligence. Data is a single artifact, information is a collection of them, and intelligence is a combination of those things, combined with context, predictive analysis, and a recommendation on how to proceed.

The GroupSense Managed Security Intelligence Center (MSIC) was born from these early lessons, and we have since made a name for ourselves in the cybersecurity industry for driving outcomes with cyber intelligence, also sometimes referred to as *digital risk*. Today, our MSIC recruits, trains, and hires some of the top cyber espionage practitioners on the planet. They speak more than a dozen languages natively and manage more than 4,000 avatars and personas across thousands of dark net, social, and chat channels, as well as private groups—even the metaverse. Graduates of MSIC have gone on to run cyber, CTI, and response programs at companies like eBay, American Express, Trend Micro, Xerox, and Google.

Cyberattackers are most successful when they have illicit access to some information, data, systems, or personnel. These digital artifacts sometimes surface because some threat actor gained unauthorized access and shared it illicitly. In other cases, though, it is simply because the rank-and-file staff have done something not malicious but silly. A mature security program utilizes digital risk and CTI to find those digital artifacts surfacing in places they shouldn't be and to mitigate the risk. This is often done by removing the artifact altogether, negotiating for its removal, or paying for its removal. In other cases, it is accomplished by implementing changes within the organization, such as new processes, procedures, or fraud playbooks, to get in front of the potential threat.

To do this well, one has to be "invited to the party," so to speak. The "party" is where threat actors, organized criminals, and nation-state agents do their business and collaborate. The "party" is often ephemeral and difficult to pin down, and doing so typically requires a massive investment in HUMINT operations. Researchers trying to find where the real dirt is being done are multilingual, context-aware espionage experts. They employ a combination of cyber know-how, personal networking, and psychology to achieve their goals.

Versions of HUMINT practices have been used since ancient times. Examples of using human intelligence in warfare date back to the Egyptians, Greeks, and Romans, all of whose armies used spies and informants to gather information. Similarly, the Byzantine Empire famously had a large network of spies to monitor its interests, and even Sun Tzu's famous *The Art of War* emphasizes the importance of spies and human intelligence sources in war. It wasn't until after World War I that modern nation-states began formalizing this practice. The United States formed the Office of Strategic Services, the predecessor to the CIA. Around the same time, the Soviet Union began operating the *Komitet Gosudarstvennoy Bezopasnosti* (KGB), or Committee for State Security. After the collapse of the Soviet Union in the early 1990s, the KGB transformed into the *Federal'naya Sluzhba Kontr-razvedky* (FSK), or Federal Counterintelligence Service, which begat the current *Federal'naya Sluzhba Bezopasnosti* (FSB), or Federal Security Service. Perhaps the Russian shell game or rebranding intelligence institutions is part of the strategy.

Russia's use of intelligence for digital espionage and cyber warfare evolved quickly. We now know, from defectors and our own intelligence investigations, that Russia's Glavnoye Razvedyvatelnoye Upravlenie (GRU), which translates to "Chief Intelligence Office," has been responsible for some of the world's most clever and devastating cyberattacks. The GRU is thought to be responsible for the Democratic National Committee breach, the misinformation campaign impacting the 2016 Presidential Election, and NotPetya, the highly sophisticated worm that caused tens of billions of dollars in economic damage. Also known as the Russian Main Directorate of the General Staff of the Armed Forces, the GRU continues to innovate and wage cyber war overtly and by way of proxy.

It is believed that in addition to the plausible deniability the ransomware operations afford Russia's Putin, they serve another purpose. Their access, skills, tools, and data capture are made available to the GRU and, thus, the Russian Federation.

The United States and other Western nations put a much stronger emphasis on HUMINT operations after the terrorist attacks on September 11, 2001. The enemy was no longer a nation-state but a

tribal, patently informal adversary operating in areas that technical tools like SIGINT or satellite surveillance could not always penetrate. Military operations in Afghanistan, Iraq, and Syria required HUMINT—human resources on the ground.

From a military perspective, HUMINT consistently provided a depth of understanding that many of the technical tools could not. While the technical tools provide vast amounts of data, HUMINT provides crucial context and human intent, including the inner plans of key adversarial decision-makers. As a result, HUMINT is often used to validate the data and information collected through other sources. Just as critical, the HUMINT operations often detected and neutralized the espionage activities of the enemy. This was a key component of counterintelligence efforts on both sides.

Despite the risks inherent in recruiting, training, fielding, and retaining HUMINT assets and resources, the spy community has long relied on HUMINT to gather information close to their adversary. Spies will often manipulate a source into providing information, perhaps resorting to bribery, blackmail, or any of a huge number of common tools used for more than a century in traditional spycraft. The cyber realm is no different—just like a CIA spy would manipulate a person to get invited to a secret party at a foreign government facility simply to plant a tiny piece of surveillance technology like an audio or video recording device, cyber spies are getting invited to the secret dark web forums. Once their fake persona or "sock puppet" is in, they plant their version of a listening device: a *scraper*, which is effectively a web crawler similar to those used by leading search engines. Many of the sources are simple web content (HTML), so this method is a proven strategy.

The scraper then feeds those conversations to a platform where the information is analyzed. How that information is processed—and whether that information becomes intelligence—differs based on the solution. In recent years the industry seems to have bifurcated into software-as-a-service (SaaS) and tech-enabled services. Both solutions aim to solve similar problems to the use cases described earlier. There isn't one correct way to do this, but there are clear differences in values and outcomes.

Both solutions rely on similar methods to obtain the necessary data. Some vendors collect this data themselves in-house, whereas others buy the data wholesale from companies that have already collected it. Still, others combine the two and use other tactical methods.

The tip of the proverbial spear in intelligence data collection is sourcing. The first step is knowing what needs to be collected and from where. The answers to those questions drive the collection technology used to gather the data. Finding the answers to these questions and gaining access to those properties are often the most difficult steps. We will circle back to source identification and access.

Once the sources have been determined and the access has been secured, technology is leveraged to gather as much information from these sources as possible. Usually this is done in the form of a scraper. There are myriad technical challenges around bot detection, CAPTCHAs, and authentication that need to be overcome in order to scrape these properties. Those challenges are real, they are complicated, and they are ever-changing—an arms race, if you will. (That's right! The bad guys use bot detection and anti-crawler technologies to keep intelligence operatives at bay. The threat actors even implement commercial solutions like Cloudflare to protect their dark web fiefdoms.) Other tools leveraging application programming interfaces (APIs) and harvesters are used to gather information from non-web properties, like chat platforms and social media.

Wholesale data acquisition presents its own challenges. There are a handful of providers that collect and sell intelligence data from the dark web and similar bad operator places. Because their model is a volume game, they collect enough data from enough sources to remain credible in the most efficient way possible. The result is a digital pipe of raw, scraped HTML data, and JavaScript Object Notation (JSON) syntax from basic sources. The more difficult-to-access sources, like sites requiring validation, specialized language skills, referrals or personal relationships, and complex tools to bypass bot detection and CAPTCHAs, are often not covered by these solutions. Unfortunately for the buyer, that is where the "good stuff" resides.

After overcoming the locating sources, access to the sources, and their defenses, another technical challenge is presented: storage and retrieval of that information. This is not a book on data architectures, but I can tell you that many companies fail at this critical component of the technology stack. The scale of what we are attempting to ingest is massive. For example, Google indexes approximately 4 percent of Internet content; the other 96 percent is deep and dark web content. Although we should not be aiming to siphon in all of that data, we should certainly target a meaningful percentage of it. Data science and architecture engineers are adept at anticipating scale and designing the necessary underpinnings to ensure a refactor or rebuild of that architecture is mitigated long term.

Further complexity is introduced when trying to leverage AI models across the data stack. Large language models (LLMs) require significant metadata structures to support the AI operation. It can be a heavy lift to implement the necessary labels, indexes, and metadata on an existing large dataset. A number of companies, like View Systems, are beginning to ease this transition for enterprises and software vendors.

Assuming one has solved for the collection of the relevant data, the next step is to make that data digestible by a human analyst. Most security platforms present data line by line in alert format, but this doesn't jive with how the human brain receives information naturally. Plus, most companies understand the need for this data but do not understand *their* why. As a result, the market is saturated with SaaS cyber intelligence companies that are essentially asking their customers questions they cannot possibly answer and displaying the results in a practically unreadable format.

The SaaS approach is usually a complex web interface front end to the collected dataset. These are usually designed to require that the customer input their own data queries on the platforms. These tools tend to be more tactical in nature and are often used in an investigatory nature rather than as a prevention mechanism. As I alluded to previously, companies often (1) don't know what questions to ask of the data and (2) don't know what to do with the answers they get. Part of this is due to the lack of experience in the commercial/enterprise

space with the intelligence process itself. Since the SaaS approach relies on the customer to run the intelligence process themselves, understanding how to develop prioritized intelligence requirements (PIRs) and how to implement those PIRs at a software level is key. Unfortunately, intelligence talent is scarce and expensive. Therefore, many companies who purchase these tools have a difficult time realizing the value of the investment. This leads to companies buying the product only to refuse to renew the software license when the first contract renewal presents itself. The SaaS vendor's primary response to this has been to offer analyst "credits" or "points" with the subscription to augment the SaaS tool with some experienced practitioners. This has worked in some cases but still falls short of running a proper intelligence program.

An effective cyber intelligence program follows the traditional *intelligence cycle* used by traditional intelligence agencies around the world (see Figure 1.1). It consists of

1. Planning and direction
2. Collection
3. Processing
4. Analysis and production
5. Dissemination

These steps are called a "cycle" because they should function as a continuous feedback loop. Most organizations purchasing a product in the cyber intelligence space benefit from the vendor's collection and often simply disseminate that data. They do not know about or intentionally skip the other steps in the process, which is the primary reason these products fail inside the enterprise.

John Holland of IntL8 is an industry expert with years of experience assisting organizations in implementing, fixing, or measuring the value of their intelligence programs. John often writes about the three tenets of a good intelligence program:

1. Cyber intelligence should be a support function.
2. You must understand your stakeholders' needs and whys.
3. It has to be relevant to the consumer of the intelligence.

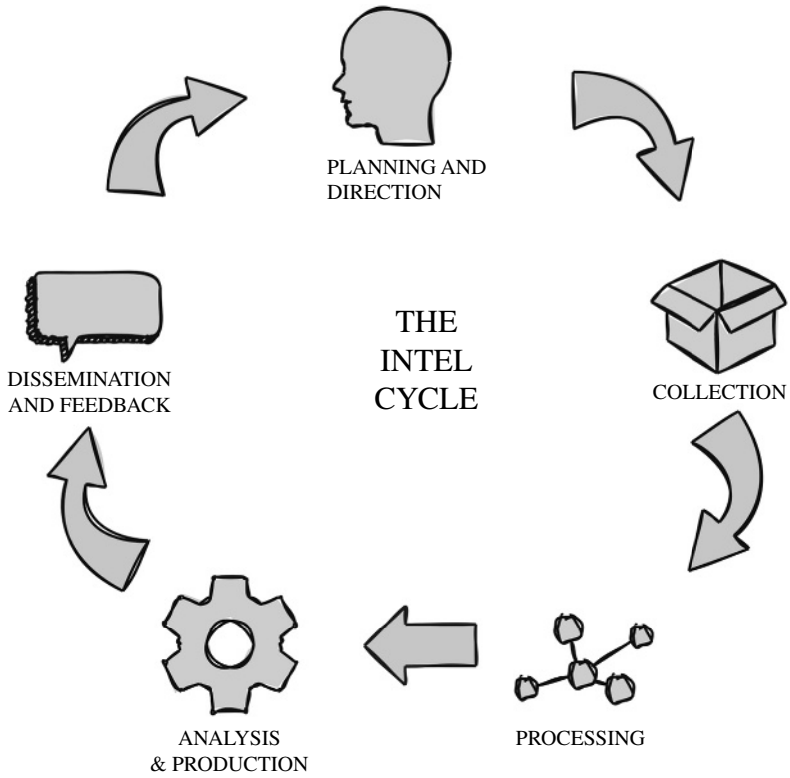

Figure 1.1 Intelligence cycle.

John's firsthand experience repairing these initiatives has shown that skipping the first step in the intelligence cycle can negate or miss the three tenets entirely for an organization.

The first step, planning and direction, is arguably the most important part of the cycle. This step is where the prioritized intel requirements are created. These PIRs should drive the next step, collection. Typical organizational PIRs should drive collection because they involve their company brand assets, infrastructure and software stack, intellectual property, projects, and even their executives or other VIPs. Cyber intelligence vendors are laser-focused on the technical aspects of collection, and those technical indicators drive their collection choices. Cyber intel vendors also try to normalize their collection to support their entire customer base.

This can cause blind spots for customers with specialized needs or strong vertical-focused risks. For example, if you are Exxon and you buy an off-the-shelf threat intel solution, that solution will certainly inform you of someone on the dark web who purports to have access to your network, but it will miss someone on an eco-warrior channel or forum who plans to blow up a pipeline. This is the purpose and the power of PIRs.

In 2016, GroupSense provided our solution as a white-label product for FireEye, a larger cybersecurity vendor. At that time, if you purchased FireEye's custom intelligence solution, it was delivered by GroupSense in the background. Standing on FireEye's shoulders, we were ushered into the largest private and public sector organizations in the world. We often discovered these organizations were upset with their current provider and couldn't understand why the intelligence solution wasn't "working." Our team, wearing FireEye shirts of course, would immediately refer to the cycle. Often it wasn't that the cyber intelligence supplier was actively or passively wrong; the organizations often didn't know what to ask of the data and didn't know what to do with it when it was received. Once enlightened, the organization gladly adopted our solution with FireEye's branding on it.

Our ability to educate the customer and understand their specific needs directly impacted how we onboard clients, as well as the technology and processes we use for collection. We quickly became one of the fastest and most flexible intel providers related to collection. When we got our first big pharmaceutical clients, we began ingesting data from sources that dealt with counterfeit drugs. When we got our first high-tech manufacturing client, we infiltrated markets that dealt with intellectual property trading and theft. We signed a cruise ship company and began collecting data related to maritime threats and information. Because our CTO, Adam Bregenzer, had experience with big data stemming from his time at Venmo, he was able to design a data architecture that became exponentially more valuable as we onboarded new clients and added collections. Every new customer benefited from the PIR-driven collection strategy from the customers that came before them.

Once data collection is tailored to meet the needs of the developed PIRs, the intelligence program must process the data. Processing consists of organizing, normalizing, deduplicating, correlating, and cleaning up the collected data. As mentioned earlier, most ingestion technologies simply scrape the raw code and throw it into an elastic search database. This results in messy results when querying the data, and it lacks any deduplication or prioritization of the data. This is why it is important to ingest the data in an intelligent way and preserve its context. Doing so is technically difficult to achieve and a bit expensive to implement, which is why many solutions don't bother to do it. If accomplished, though, it makes the next steps in the intel cycle go more smoothly.

Many vendors in the cyber intelligence space skip or cheat the analysis and production part of the cycle. Similarly to contextual ingestion, this can be technically difficult and expensive to do well. It is important, nonetheless, as this is the step that *transforms data into intelligence*. Considering that this is not something many of the vendors provide, the product the buyer receives ends up being simply *data or information* rather than true *intelligence*. The vendors, perhaps unintentionally, market their wares as intelligence tools effective at producing . . . intelligence. This leaves the onus on the consumer of the data to produce the intelligence. Unfortunately, most enterprises are not staffed or trained to do this well, if at all. The result is a consumer upset with the value of their investment in the tool, even if in an uninformed way.

The analysis function is best performed by a human trained in the intelligence discipline. The analyst on the receiving end of the data has to reformat that data in a way the customer, whether an individual, a department, or an entire institution, can understand. The analyst has to add a summary of the impact and who or what is affected, contextualize the data, provide some expected outcomes, and recommend remediation action; only then is intelligence produced. At GroupSense, we recognized early on that the total addressable market for data that needed to be turned into intelligence was quite small. Most enterprises did not have the staff with adequate talent, the resources, or the necessary processes to

take data from a platform and create intelligence. This is precisely why, in our early days, we began recruiting and training cyber intelligence analysts to "bookend" our platform known as Tracelight. The bookend, as the label suggests, means that our analysts assist the customer in creating the right questions to ask of the data in the platform. The system does the heavy lifting by ingesting data from sources, cleaning and deduplicating it, correlating everything, and adding any additional technical context for the analyst. This creates information from the data, which is then fed to the analyst. The second bookend is the analyst taking this information and creating intelligence for the customer. Only then is an advisory or report published inside the platform. This approach has worked well for the 11+ years we have been building on it. I am encouraged by and proud of the difference we have made for our clients.

The last step, dissemination, is simply making sure the intelligence gets into the right hands quickly and in an easily consumable manner. When possible, integrations with other platforms, such as ticketing systems, dashboards, business intelligence software, and antifraud technology, ease the consumption of the data for the customer.

It is just as important that it does not end up in the wrong hands. The technology and tradecraft, sources, personas, sock puppets, and activities related to the intel process are critically secret. Compromising any of these can create opportunities for counterintelligence or render the whole operation useless. This is why companies like GroupSense hold these functions, indicators, and data so closely. This is also why GroupSense does not buy collection data from third parties. It is crucial that we control the process closely to ensure the data that we turn into intelligence is legitimate. We have been awarded patents on how we accomplish this.

Of course, doing any of this can be largely useless if the intelligence is not relevant or actionable to the consumer.

In Chapter 3, you'll learn how to find and infiltrate sources using personas and sock puppets. It takes careful steps and sometimes years to develop personas that can be tasked effectively in the underground market. This is why many companies, aside from the

legal risk, lean on their intelligence provider to take action or validate the intelligence they receive.

If the intelligence suggests that the threat actor has access to internal systems, the intelligence advisory will likely provide context and some dossier on the threat actor making those claims. Does this threat actor normally sell this kind of access? Have other threat actors had successful transactions with this person? How active is this actor on this or other illicit forums? What number of transactions has this actor performed? This information would be followed by a suggestion or action. For example, due to the threat actor's relatively low transaction history, the intelligence advisory may suggest engagement with that actor to get some kind of proof of the stolen access. If the customer approves the engagement, a researcher from our MSIC team will select a persona that makes sense for the engagement, relying on one regularly engaged in this kind of access acquisition or interest in it. The persona would have to be a tenured entity in this particular forum. Other threat actors must say good things or validate this persona as a known bad actor; vouching for them in this environment is critical.

Once the persona is selected, the researcher from MSIC will engage the threat actor. More than likely, the conversation will start on a forum as a reply to the advertising thread. It will likely be in the native language of the threat actor making the advertisement, which is often not English. (Perhaps Russian or Ukrainian; of course, the MSIC researcher who is most fluent in that language would have been assigned.) The conversation typically transcends to a direct chat utilizing another platform, perhaps TOX or Jabber. At this point the persona will indicate interest and will ask the threat actor for some kind of proof. Often the threat actors will provide this proof in the form of a screenshot or a system command-line command. That data, after running through a sandbox and checking EXIF data (metadata inside any images that might hold key information about our alleged attacker), is provided as an update to the original intelligence advisory. Again, if there is a logical action to be taken, it will be provided as part of the advisory.

Another example of actionable cyber intelligence occurred with one of our large international pharmaceutical clients. We provided them table-stakes intelligence around cyberthreats on the dark web, Telegram, and other illicit sources. In cases where threat actors signaled their possession of inside information or access to the company's systems, we provided the finished intelligence to their IT security team. Each intelligence advisory we provide includes recommended next steps or actions to be taken. In many cases, those actions will be taken by our MSIC team rather than the client. Generally, enterprise companies are not willing to take on the risk of interacting with threat actors directly, nor do they have the OPSEC infrastructure or capabilities to do so. In some cases, purchases may be made on behalf of the client, which requires lawyers, contracts, financial agreements, and sometimes an affidavit to ensure compliance. These purchases are typically done through cryptocurrency and often involve a marketplace-supported escrow service.

In the case of the pharmaceutical company, we issued an advisory about a Russian actor who purported to have quantities of a rare and carefully regulated drug produced by the pharmaceutical company. Naturally, the pharmaceutical company was concerned. They were particularly interested in the presence of an ingredient in that drug protected as a trade secret. Curious about how that drug ended up on an underground market, they asked MSIC to engage the threat actor to gain more information. The MSIC team had a few active Russian-speaking personas in the counterfeit drug marketplaces, chose one, and reached out to the actor. After much back and forth, it was determined that it was likely stolen material. Our team arranged, through an intermediary, to have someone purchase this counterfeit item from the actor. In this case, the actor had worked out a deal with a brick-and-mortar pharmacy. Our agent, whom I call "Russian Todd" because he looks like my friend Todd, exchanged cash in person for a brown bag of medicine. Russian Todd then sent the medicine to a lab owned by the pharmaceutical company. Although we didn't get a final report from the client, as it was their private business, we assumed the threat actor was working for a production facility near the pharmacy—an inside job.

Needless to say, GroupSense is likely one of the very few providers in the cyber intelligence and digital risk space currently covering the intelligence cycle entirely and assisting with kinetic outcomes when necessary.

Fairly early in our journey, we were approached by one of our large municipal clients. The client was one of the largest U.S. cities and had a very outspoken mayor. The TI team at the city's cyber command division asked if we could take what we were doing for the city itself and apply the methods to protect the mayor. We thought that theoretically, while maybe some of the sources would be different, reducing a person's digital risk is a very similar use case and should work. Through this experiment, we developed our VIPRecon product that uses a tech-enabled, intel cycle–focused service approach to protecting government officials, VIPs, high-net-worth individuals, and celebrities. Today, GroupSense covers the gamut of digital protection for *people* as well as organizations.

This is how cyber intelligence solutions should work, by developing PIRs based on business or personal requirements that *drive* collection requirements. The data collected gets processed by computers or humans. The remaining information is contextualized and paired with impact assessments and recommended courses of action. Then actions, either offensive or defensive in nature, are taken. That is *intelligence*, friends.

A true cyber intelligence provider is a fundamental part of the customer organization's cyber team and overall cyber program. The provider should work closely with the leaders, cyber and otherwise, and with the boots on the ground to ensure that the necessary actions are taken, risk mitigation strategies are implemented, and results are fed back into the intelligence cycle to improve the overall program—in other words, provide outcomes, not *data*.

In August 2024, Intel471, a provider of cyberthreat intelligence, sponsored a partnership of 28 other industry participants to form the Cyber Threat Intelligence Capability Maturity Model (CTI-CMM; `https://cti-cmm.org`). The CTI-CMM aims to support cyber intelligence practitioners in the field by providing a framework for cyber intelligence programs. Some of the brightest minds

in the industry are contributing to this project. I am encouraged that progress is being made.

If you do one thing after reading this chapter, please familiarize yourself with the intelligence cycle. If you are a business leader, try to understand what elements of the cycle are being missed today and what kind of solution your provider is offering. If it is simply data or information, do you have the staff and intelligence talent to turn that information or data into intelligence? Most importantly, when you have finished crafting the intelligence, can you do anything with it?

Cyber Recon Leader Profile:

Brye Ravattine

Bryeanne Ravattine

Current Role: Director of State, Local, and Education for ShadowDragon

Other Affiliations: Liberty University Athletics Hall of Fame Inductee, Craig Beardsley Award Winner, Presidential Management Fellow finalist

Brye's Bio

Bryeanne "Brye" Ravattine is a cyberthreat intelligence powerhouse passionate about solving tough problems and making the digital world safer. As the Director of State and Local at ShadowDragon, she helps law enforcement and government partners track down criminals and protect their communities using cutting-edge open source intelligence tools. Before diving into cyber, Brye competed at the Olympic Trials and is an NCAA D1 swimmer, where she mastered the art of focus, discipline, and pushing past limits, skills she now brings to every mission in cyber.

When she's not working cases or speaking at conferences, you'll probably find her hanging out with her beloved pups, hiking somewhere beautiful, or catching up on her favorite true crime podcasts. Brye's a big believer in blending strength and empathy, and anyone who's worked with her knows: she's as sharp as she is kind.

Brye's Mission

The problems we face in the world, especially in the digital space, can feel overwhelming and sometimes impossible to solve. But I believe that every action matters. Every small step forward makes

it harder for those who seek to harm, and I've dedicated myself to being part of that resistance.

I come from a family rooted in service—military, law enforcement, and a deep sense of duty—and that legacy fuels me daily. My moral compass keeps me grounded, and my drive comes from a simple but powerful belief: If you can help, you should. Whether it's supporting a community under threat or standing up for someone who can't do it alone, I'm here to make a difference—no matter how big or small.

I'm a relentless do-gooder at heart, and I find real fulfillment in lifting others up, fighting injustice, and using my skills to protect the vulnerable. That's my mission, and I'm just getting started.

YouTube URL: https://youtu.be/8TgThakROnE?si=eB0SRQYwPT7voKM9