Disrupting Nation State Hackers

Rob Joyce
Chief of Tailored Access Operations
National Security Agency
@ USENIX Enigma 2016

Appreciate it, thanks for the welcome. So, as David introduced, I'm from Tailored Access Operations. And I will admit, it is very strange, right, to be in that position up here on a stage in front of a group of people – it's not something often done. But, I'm in a unique position in that we produce, in TAO, foreign intelligence for a wide range of missions, to include advice to informing policy makers, protecting the nation's war fighters 24/7, and in that space, we're doing nation-state exploitation. And so, my talk today is to tell you, as a nation-state exploiter, what can you do to defend yourself to make my life hard, right? So, not many people will stand on the stage and have the perspective of an organization that does exploitation and to be able to talk to those elements that really would disrupt the nation-state hackers. So, in that vein, I want you to think about if there's something you really, really want to protect, what do you have to do?

So, you'll hear a common theme throughout my talk. It'll boil down to a couple small things. The theme I want you to take away is, if you really want to protect your network, you really have to know your network. You have to know the devices, the security technologies, and the things inside it. So why are we successful? We put the time in to know that network. We put the time in to know it better than the people who designed it and the people who are securing it. And that's the bottom line. And you'll kind of hear that woven throughout the talk. So, if you think about what goes into an intrusion, there's a series of phases that happen, right? As you walk down through these, I'll talk about the things that can—that we focus on. And you could break the chain throughout that compromise by disrupting the transitions between these elements.

So, really, the first phase during a targeted intrusion is a reconnaissance phase. Somebody's got to go out and understand the target. It starts with simple things, like scanning. Go out and physically scan the actual target. There's understanding important people, or email addresses from that activity. Going out and looking at the open-source information about that target. So, it really is, what can you learn? What can you understand? As I said earlier, our key to success is knowing that network better than the people who set it up. So, in that space, the reconnaissance phase is really important. I'm gonna move my laptop a little here, so I can get to my notes.

So another key point inside this, you know the technologies you intended to use in that network. We know the technologies that are actually in use in that network. Subtle difference. Did you catch that? You know what you intended to use; we know what's actually in use inside there. So, when we look at that, we will learn the security functionality of the devices inside that network, we'll study them, understand them, find the vulnerabilities. In fact, we've got people who will understand the security functionalities of those devices, better than the people who developed the actual device, right? So they won't know the whole product, they won't know every feature that the developers had, but they'll understand the security technologies, and they'll bring that expertise at a very, very deep level. So inside that, it's minute attention to detail inside that security layer, again, knowing the network, knowing that space. So what does that mean? We apply the focus and energy to look at those details. Will you, as people who have important things to protect and hold dear, will you put in the energy to understand the network, understand the devices, and configure and use them in the proper way that would prevent exploitation? So, there's a foundational piece of advice to countering these kinds of threats, right? You've got to have procedures to evaluate what you'll use, what you'll install. You've got to lock down and disable those things that you're not using, right? Reduce the attack surface. It's not a new or amazingly insightful piece of advice. But you'd be surprised, as I said, about the things that are running on a network, versus the things that you think are supposed to be there.

So what can you do to understand that exposure surface? Red team that network. Bring in pen testers. Poke and prod it, just like an adversary will do, to find out what's inside that space. Find out what's exploitable. Well-run networks really do make our job hard. So, if you go to the trouble of understanding what's inside a network, you run that pen test, you've got those results, act on it. So NSA, in our information assurance side, will do red team testing against government networks. So we'll inevitably find things that are misconfigured, things that shouldn't be set up inside that network—holes and flaws—and we'll produce reports telling the network owner things they need to fix. Cycle comes around to the point where we've got to get back and redo a red team against that same network. It is not uncommon for us to find the same security flaws that were in that original report. That's the first place we go, is to the original report. Did the things we pointed out previously get fixed? So, inexcusable, inconceivable, but returning a couple years later, the same holes and vulnerabilities exist. I've seen it in the corporate sector, too. I've seen it in our targets, right? People tell you you're vulnerable in a space, close it down and lock it down. So if you've invested the resources, to do that kind of discovery and red team space, go ahead and follow through.

Another key point, don't assume a crack is too small to be noticed or too small to be exploited. So if you go through and do that pen test and you say, "We look great on these 97 things, but these three things over here, they're kinda esoteric. They probably don't matter much. We'll probably ignore them," right? That's what we need. We need that toe hold. We need that first crack, that first seam, and we're gonna look and look and look for that esoteric kind of edge case to break it open and to crack in. So pay attention to those results. Same thing in this discussion about the temporary security vulnerabilities. So if you own a network, and you've got trouble with an appliance inside your trust zone, inside your network boundary, and you're talking to the vendor, and just can't quite make it work. And they say, "Well, open it up for me. I'll come in. We'll poke around. We'll take some logs. We'll fix it for you. We'll do it over the weekend, don't worry," right? Are you gonna open that door for that 24, 36 hours? So I'll tell you…the nation-state attackers? There's a reason it's called advanced persistent threats. Because we'll poke and we'll poke and we'll wait and we'll wait and we'll wait, right? We're looking for that opportunity—that opening and that opportunity, to finish the mission.

Another big area, I'd say, in this reconnaissance phase, is figuring out about the network boundaries. So, I talked earlier about you know the things you intend to have in your network. We look for the things that are actually in your network. Well, that's becoming harder and harder these days as the network boundary gets more amorphous, gets more porous, or gets more inclusive of other things. Think about trends like bring your own devices. Internet of things, work from home access. These have really created situations where internet—interconnected network elements are under varying administration control, right? I even see the case where leased facilities come with a leased network that is under the control of that physical location, and trusted in interne—interconnected to your domain, right? So, think about the things that are now a component of your domain, your trust zone. Cloud computing, right? Cloud computing is really a fancy name for somebody else's computer. If you have your data in the Cloud, right, you're trusting the security protocols, the physical security, all of the other elements of trust in an outside entity. May be done right, it may not. You may have varying degrees of understanding about what's inside that cloud. But they are now part of your risk and liability. So, I see a growing trend that are really making it hard and diffusing the network boundary. Trust boundaries now extended to partners, personal devices, right? All of us love to have our iPhones, Androids, tablets, devices come and go, right? You're trusting those onto the network. There's even the heating and cooling systems, right? Other elements of building infrastructure and more. So what are you doing to really shore up the trust boundary around the things you absolutely must defend? And that, for me, is what it comes down to. Do you really know what the keys to the kingdom are that you must defend, right? Instrument, defend, pay attention to those crown jewels, because that attention and rigor really makes our job hard.

So after reconnaissance, the next phase is gettin' that initial exploitation. Gotta find a way to get energy inside that network, can you go ahead and get some opportunity. These things can happen from spear fishing, they can happen from water-holing, is there a weekly defendant site that everybody goes to? Exploiting a known CVE, right? There's already a vulnerability and there's a recipe for exploiting that activity, already done. SQL injection. Exploiting a zero day, other technologies, ways to get in. I think a lot of people think, the nation-states, they're running on this engine of zero days, you go out with your master skeleton key and unlock the door and you're in. It's not that. Take these big corporate networks, these large networks, any large network – I will tell you that persistence and focus will get you in, will achieve that exploitation without the zero days. There's so many more vectors that are easier, less risky, and quite often more productive than going down that route. So, to ward off a persistent actor, you really need to invest in continuous defensive work, right? Because if the CVE world is continuously rollin' and pumpin' out new information about cracks and holes in existing products and services, you've got to be continually updating and defending inside that space. So most intrusions come down to one of three initial vectors, right? Email, where a user opened an email, clicked on something that they shouldn't have. A website, where they've gotten to a malicious website and they've gone ahead and it's either executed, or they've run content from that website, or removable media where a user inserted contaminated media. Sometimes even bridging an air gap network, right? But those three are the big three. Where do you need to go in this space? You really need to get the networks not to rely on the users to automatically make the right decisions. Sometimes, even the experts get it wrong.

So how can we build and insure the policies and the technical enforcement of those written policies – keep accidents and slip-ups from occurring, right? Because, I don't care how many times you train people from not clicking on those unsolicited emails, people do. And even when you get to the nation-state advanced persistent level, sometimes those emails can be really well-crafted, to the point where it's not an unreasonable thing for somebody to click on. So how do you prevent that from detonating? Can your architecture and your policies defend against those user actions that are gonna take place? Can they stop those threat factors? Because if they can, it really makes my job hard. So, one thing I'd absolutely recommend is things like anti-exploitation features. Microsoft EMET, everyone ought to be turning that on. It really does slow down the amount of vectors that are available for something to execute in that space. So, I'd look at NSA information assurance directorates, they have a host-mitigation package, so it's best practices for locking down and mitigating at the host level. EMET is only one of those recommendations – there's a whole series of things that really do lock things down well. That's the guide, those are the specificity, there's not the secret sauce that goes beyond that, inside the protection of classified material for the US government, right? Look at that guide; it really, really is solid.

The other thing you've gotta do, you've gotta take care of—take advantage of software improvement. I mentioned CVE's and vulnerabilities. Boy, if there's a known bug in a software that's exploitable, you ought to be fixin' that and getting it off your network. So, I think, tip of the hat to the software industry that is making upgrades and automatic patching a background activity that's beyond the user control, right? That is an outstanding security practice, where it is just taking care of every time there's a newly closed vulnerability, it becomes part of your ecosystem. That's an outstanding thing and that cuts down the opportunity window between known vulnerability and execution. And if that patch window is months or years, again, an inexcusable practice. So the other thing I'd encourage is, use a secure host baseline. So, again, that kind of goes like the host mitigation plan, the IED product. Secure host baseline is the current best practices for locking down configurations. Again, there's some out on the NSA information assurance website to look at. So I'll tell you, our organization teaches and trains. That's one thing we do really, really well, right? We institutionalize that knowledge, we teach people to get them to the next level so that they can work and exploit. So we train best practices. We pass those on, we use those best practices. So I'm gonna use best practices for exploitation. Are you gonna use best practices for defense? Again, it really comes down to that. If you have something somebody's coming at, and you need to defend it, you need to be looking at, what is that apex predator gonna be doing to come after your information? They're gonna be using best practices for offense, you've gotta be using best practices for defense. In almost any intrusion, you know, in this initial exploitation space, people are trying to get credentials, right? Often legitimate credentials are compromised, enabling intruders to get in, and masquerade as legitimate users, coming after the network. And it's imperative that you have some processes and plans to understand what normal is inside your network.

So, if somebody’s got credentials, are they operating under the norms for those credentials? Are they going places that they should be? Are they trying things that they shouldn’t be doing, right? Better-defended networks require specific methods for accessing the resources of that network. They monitor credential uses, they look for anomalous behaviors, two factor authentication, right? Making it that much harder to steal credentials, and it really is important to make sure that that small crack of a lost credential doesn’t get turned into a pivot in a later stage into a large access. There’s been numerous security best practices that have been recommended over the years. But, some of the things, like, making sure least privileges for accounts. There are only a very small handful of accounts that have the keys to the kingdom, and you only give the privileges needed to specific users. Not everybody’s happy livin’ in that world, right? Why can’t I have admin into my server, my boxes, those kind of pieces, those are the kind of wide-ranging credential reuses that wind up turning into large-scale compromises. Segmenting off portions of the networks, rarely implemented, wait-listing, things like that. If you care about your things, consider those. They really do make our life hard. We also really love it when administrator credentials, or other system-wide credentials are hard-coded into scripts, or accessible on the devices.

So, I think people are starting to understand the pass-the-hash vulnerability, right? If you haven’t learned about that, if you don’t know what pass-the-hash is, go understand it. So, that’s something where you can get a domain credential and, you can grab a credential and move laterally onto other machines and just pivot like mad throughout the network. So, one of the key activities is really thinking about how you manage those capabilities so that you can protect against pass-the-hash. I mentioned that if things are hard-coded and included in scripts, you know, they're vulnerable and likely to be pulled. Most of the modern protocols these days are not passing credentials in the clear, but do you think nation-states are taking advantage of the ones that are? Right? So you gotta look for those older protocols, drive them out of your networks. It's not enough to know about things like pass-the-hash and making sure that all of the authentications are done only with more modern protocols that keep the passcodes and passwords out of plain text, but think about where you've hard-coded and enabled one box to log you in through an account to another to do an activity. It really does make yourself vulnerable. The other big thing I'd recommend, enable those logs, but also look at the logs. You'd be amazed at incident response teams go in and, you know, there's been some tremendous breach. "Yep, there it is right there in the logs." Great, you've got logs, it'll tell you that you've been had. Enable those logs. Look at those logs. I'll tell you one of our worst nightmares is that out-of-band network tap that really is capturing all the data understanding anomalous behavior going on and somebody's paying attention to it. So rewind all the way back to the beginning of my talk where I said, "You've got to know your network, understand your network, because we're going to," right? Those logs, they are just the rock-bottom bedrock foundation of understanding if you've got a problem, or if you've got somebody rattling the doorknobs to give you a problem. Right?

So somebody's cracked open the door, they're on the threshold. The next thing they wanna do, is they've gotta establish persistence. It's not good enough just to be in a network, but if you're really there to exploit you want to dig in and hold, right? So work happens at this point. Privilege escalate maybe? So that you can get down some tools. Finding run keys, getting into scripts, other technologies to ensure that persistence onto those computers so that you can stay. One of the things we run into here, things that have implemented application white listing, makes this world hard. Application white listing, it is difficult for generic users in a large network to know exactly what applications you're gonna run, what should be permitted. There's some good work going on to make this a little more generic and understand what's routine and what's not inside and organization. But, again, like I said, figure out early what you need to protect, segment that off, and that's the place you maybe want to think about whitelisting, right? Make sure that in that space they can't run a piece of malware, or something new or unusual. Your goal needs to be to main—to restrain that malicious behavior. Keep it from launching in the interim.

So then, after you've gotten into the network, install some tools, right? Usually the first tools down are lightweight, small beaconing things. Their intent is to establish that beachhead and then bring down the tools that are actually gonna do the work. Um, so, so, there are things I think the AV industry at times gets a bad rap for their ability or inability to keep things off, you know. If your AV is a list of bad things that shouldn't run on your computer, that's not a great technique because that just means the unique thing you need to run on that computer needs to be unique and it will never be in that list. But the research and the technology's evolving now where reputation services are more the norm. So, every piece of software that wants to execute on your machine, gets hashed, pushed up into the Cloud. Let me tell ya, if you've got a reputation service and it says that interesting executable that you think you want to run in the entire history of the internet has been run one time and it's on your machine? Be afraid, right? Be very afraid.

So, reputation services are a growing technology that can make our life hard. Similarly, most of these tools want to talk out to a domain to get those further modules. They want to talk out and call back home. They want to report success, or bring data back. So, so they'll be wearing a domain name, right? Reputation services work probably even better in the domain name world, because the domain names, if—it's not enough to block bad, known bad domains, right? That's important, but usually that'll get you the crimeware. You've gotta block the things that are not known good. It's really hard for an exploiter to get a website created and established that has good reputation. It's not hard to register a domain and make something call out to it. But, but if something is evaluating that information, and nobody else is going to it, or the content's stale, it's not updated, it will have neutral or negative domain— or neutral or negative reputation. So, again, reputation services looking at that, that's a hard thing to overcome in domain names.

So, after you're in a network, rarely do you land where you need to be. At this point, it's important to move laterally and find the things you need to find. So, the big question you need to think about is, if you have an intrusion somewhere in your network, can you then defend against this lateral movement? If you think about it, most networks: big castle walls, hard crunchy outer shell, soft gooey center. How do you get to the point where you know you have an intrusion and you're gonna keep somebody and make it difficult for them to move from the place they landed to the place they need to be? And so, again network segmentation, monitoring, caring about your – the accesses that allow these privileges. They're all really important pieces. So advanced attackers really go for the crown jewels, right? They're gonna go for those domain admins to control the entire network. You really need to limit the administrator privileges. Segment the accesses, enforce two factor authentication. Nothing is really more frustrating to us than to be inside a network, know where the thing is you need to go get to, and not have a path to get over to find that.

So the other thing is, you know, poorly considered trust relationships. I talked earlier about the amorphous edge of your networks, allowing any network— any user or any net computer with valid credentials to access the network from anywhere, that's a poor idea. A huge risk. Better networks employ things like comply to connect for remote access. They connect and assure the security of the remote connections. Maybe even figuring out physical locations, where you're calling from in. Seen some really interesting things with dynamic privileges, thinking about, you can access pieces of information from inside your network, but not from out; inside the state, but not out. So, there's ways to limit and consider the segmentation in a creative way. If you really wanna make my life hard, you segment, you manage the trust to the most important places. You consider who really needs that trust and who should be able to access those things. I think another key thought that people don't have, is consider how – consider that you're already penetrated, right? Do you have the means and methods to understand if somebody's inside your network. If you, if you read statistics, Verizon does a great intrusion report every year. Look at the statistics for how long intrusions go undetected. Months or years, right, after people are inside. So what do you have to understand and contain after that first pieces.

So monitoring and detecting inside the networks is just as important as that network boundary. And, and many networks they don't have incident responses, response plans, and if they do they rarely exercise them, right? Have you ever seen incident response plan exercised inside your network? So, the internet of things, the boundary conditions, all bringing things that are probably untrusted inside your network. Why go after the professionally administered enterprise network when people are bringing their home laptops? Their kids were going out and goin' and downloadin' Steam games the night before, right, inside your network and trust unit. What's that trust boundary?

And then, as we mentioned earlier, the internet of things, there is now getting to be a whole SCADA network running in parallel, sometimes interconnected to your whole corporate network. Have we thought about those security elements? Ron Rivest, you know, made a great point earlier today. Have we got those things right? Do we need to invest more in tho— those technologies to secure and defend there? Absolutely. So at that point, we own you. All that's left to do is to collect, exfil and exploit, right? So, once inside a network, the main focus is getting what you need, getting it out, and leaving undetected. So data theft is one arena, but I challenge you to think about a new one, right? In the wake of Sony attacks, everybody's gotta think about, right, I've got my basket of eggs, I've got my most important things. I've defended them, I've instrumented them, I've packed them ever so carefully in that bubble wrap and kept it off to the side with my best security practices…what about the destructive attack? So offsite back-ups need to be part of your plan, figuring out how you're gonna deal with data corruption, data manipulation, or data destruction. It really needs to be something you're thinking about now. Don't be that Saudi Aramco, that Sony, that learns about it afterwards and then is improving. You gotta think about it now.

So, the other thing I'd point out is, you gotta differentiate between the cyber criminals and the nation-state intruders. So, last weekend, we had the huge snowstorm on the east coast. Turns out my neighborhood, in the middle of the night, one guy walked through the neighborhood, came through the whole court, checkin' every car door to see what was unlocked. Took anything that wasn't nailed down in unlocked cars. Didn't break a window, didn't pick a lock, just took opportunistically whatever he could, right? That's a lot of the internet malware, badware. It's looking for credit cards and opportunities to use your machine to send spam and make money. To do cryptolocker and lock down and extort you for money. But at that point, you know, they're opportunistic. They're looking for the back, weak gazelle in the pack to pick off. If you're looking at the nation-state hackers, we're gonna be persistent. We're gonna keep coming and coming and coming. So, you've gotta be defending and improving and defending and improving and evaluating and improving, right? The static person is gonna float to the back of the pack and not for the crimeware, but for the nation-state advanced hacker, they're gonna find those CVE's, those things that are not patched, they're gonna find ways in that aren't monitored, they're gonna steal credentials, they're gonna get to those pieces. So don't be that easy mark.

Anybody holding up the camera? Who's gonna scan the QR code from the NSA guy? All right. So, that is a link, it's a real link, it's not a Rickroll, I promise. Trust me. So I'd encourage you to go to the NSA website, there is some awesome material that keeps you from being at the back of the herd, right? It is tough to defend against that nation-state advanced persistent threat, but you really can make a huge, huge difference. So, you ought to be tightening down and learning some of these lessons, right? So thank you for your time and attention.

[Transcript by Christian Folini]