

Red Teaming, an Interview with Ray Park of Sandia National Labs (SNL)

Listen to Podcast [here](#).

Brian Contos: Welcome to the Imperva Security podcast. I'm Brian Contos, chief security strategist for Imperva. If you would like to learn more about this subject and Imperva, visit [Imperva.com](https://www.imperva.com), check out our [blog](#), follow us on [Twitter](#), or send us an e-mail at blog@imperva.com.

Brian Contos: Joining me today is Raymond Parks, a program manager at Sandia's Information Design Assurance Red Team, or IDART, a senior member of the technical staff in the Assurance, Technology, and Assessment Department at Sandia National Laboratories, SNL, and project lead for several control system security projects.

He has led red teams through assessments and has been a team member of over three dozen other red teams. Currently, he is leading the red team assessment for the financial critical infrastructure element. Recently, he led a team that performed a North American Reliability Corporation, or NERC, Critical Infrastructure Protection cyber vulnerability assessment for a major Midwest utility. Ray Parks is a graduate of the United States Air Force Academy, with a bachelor's of science in engineering, and he is also an Eagle Scout.

Welcome to the show, Ray.

Raymond Parks: Well, hello there.

Brian Contos: Ray, before we get into our conversation about red teams, I did want to ask you, and while this podcast is going to come out a little bit after the fact, a recent report came out, April 2009, claiming that "spies" have compromised the US electric grid. And there's been a lot of press about this; I think a lot of speculation. I was wondering if you could just sort of give us your personal perspectives on this and what you think it may or may not be.

Raymond Parks: Well, I'm going to be speaking personally here, not as an employee of Sandia National Labs. But basically, I really don't know much about it beyond what's in the press. Some of us have just been sending emails back and forth saying, "Hey, look, somebody else wrote another blog about it," or "Somebody wrote another article about it." But, it's not very clear what exactly happened and to what extent it happened and why people believe that it's spies.

I said elsewhere in another forum that the biggest difficulty is not that somebody may actually have penetrated a control system computer. That's not the hard part of the process. The hard part is trying to figure out what to do when you're there. And that may be something that a lot of people have ignored is that it's not a simple thing to do, to just get in there and start turning the lights off. It doesn't work that way.

Red Teaming with Ray Parks

Brian Contos: So, Ray, let's just jump right into things, then. I've gone through your bio, but if you could give our listeners some background on what exactly it is you do, and maybe some background in how you actually got into this field as well.

Raymond Parks: Well, right now, I'm working both red teaming and critical infrastructure control systems. There's some overlap between the two, but to be honest, we don't use our red team control systems directly. My red teaming work has been with government agencies who have or are critical systems vital to national security. And really, that's a really fun kind of work, because you know you're making a difference.

Some of the recent control system work is just as exciting. We're combining a tool that we use for our red teaming to analyze attacks, which we call a Graphical Adversary Modeling Environment, or GAME, with the virtual control systems environment we built for DOE. And by hooking the two up, we're going to be able to provide operator training of what it's like to experience a cyber attack at the IFIP meeting at the end of April.

Brian Contos: How long have you been doing this type of work?

Raymond Parks: Well, I've been working on red teaming since 1996. I graduated from the Air Force Academy in 1978 and worked for about nine years on active duty as a programmer and a manager of programmers. And I got my first control system experience during that time. That's not something people normally think of as control systems, but I was responsible for the majority of the ground system that controlled a constellation of satellites.

And then, after I left active duty, I was a consultant and analyst for various Navy and Department of Energy programs, before I started working at Sandia Labs.

So, as I said, I started red teaming in 1996. I became part of our Information Design Assurance Red Team. But, I have to admit, I've been involved peripherally, because I think everybody has to be, with security long before that. And even as being a bad guy. I hacked my first government computer when I was in high school, in 1973.

Brian Contos: [laughs]

Raymond Parks: I helped resolve a time bomb that was left behind by a disgruntled employee and closed a network back door a contractor was using during my satellite days.

After I left active duty, I also worked as a reservist at the Air Force Weapons Laboratory in the Nuclear Surety Division. Primarily, what we were trying to do is figure out how some bad guy might be able to get at or use nuclear weapons, and trying to make sure it could never happen, which kind of led back into the red teaming work, which I've been doing since 1996. And about five years ago, I got back into control systems, primarily in the area of security.

Brian Contos: So, when we say the term "red team," what's that mean? Who are the people that make up red teams? Are they only government employees? Are they civilian agencies? What exactly is this?

Raymond Parks: Well, different folks have different definitions of red teams. And in fact, about 10 years into our work, in 2006, those of us that were working at IDART decided to develop some guidance for people that wanted red teaming - program managers and sponsors.

Red Teaming with Ray Parks

So, we created a course. We called it Red Teaming for Program Managers, or RT4PM. And in doing the course, we had to, of course, settle on our own definition of red teaming, because obviously we couldn't be telling somebody else about it if we didn't know what we wanted to call it or how we wanted to define it. And so we defined it as "authorized, adversary based assessment for defensive purposes."

Authorized means someone with legal control of the facility or system or entity to be red teamed has agreed to the process. You have to be authorized, otherwise you can go to jail.

Adversary based means that the activity is centered around what would one or more adversaries do if they were attacking the target? This means taking into account their knowledge, the adversary's skills, their commitment, their resources, and even sometimes their culture.

Assessment means one is making a judgment, possibly a comparison, of the state of the target with respect to the actions by the adversary. So, you're making a judgment about what the system looks like and how it's working and how secure whatever you're looking at is.

Now, saying that, I would point out that we deliberately excluded security in our definition. We didn't say "security assessment" because red teaming doesn't necessarily involve security and attacks. We have done red teaming of adversary reactions to potential business decisions, and there's a whole group of people in the government that do red teaming to try to understand how other governments and militaries are going to react to what we do.

Defensive purposes ties us back to the good guys. We do what we do to help the good guys make decisions, about business, about security, about computer systems, and about control systems in my particular case.

Brian Contos: So, when you go off on a job or a contract and you're doing some red teaming work, is it usually the organization itself that's requesting this service, or is it a government organization that's auditing them? How is it that you come into your projects? Who's making the request, generally?

Raymond Parks: Well, we prefer it if it's the owner of the organization that wants us to red team them, because then that means we're going to probably have the most success because of the cooperation. But, there have been times when we've done what we call third-party red teaming, where somebody, either a government agency or some other entity, has paid for the actual red teaming, and the folks who own the system are the recipients of the red teaming. And those can be good, if we can establish a good relationship, but sometimes they're not all that great.

Brian Contos: Interesting. Now, a lot of people think of testing of facilities that house nuclear power plants and turbines and this type of thing. They associate that with physical security attacks, and organizations trying to break in physically through the roof or through whatever. Is that something that your team does as well, or do you partner up with other groups that specialize in the physical portion of this?

Raymond Parks: We usually partner for physical security types of work. That's not our primary area, but there're other folks here at Sandia who've been doing that for a very long time, much longer than we've been doing the cyber side of things.

Red Teaming with Ray Parks

Getting back to my definition, a lot of folks are doing work that fits within our definition of red teaming. Within that RT4PM process that we came up with, there're about eight different types of red teaming, which we basically developed from our own experience and from what we learned from other folks during our series of red team conferences that we had. And most of that, from my standpoint, anyone who does any of those types of activities is a red team.

There are military red teams, which I'm sure you've heard of. The NSA and each of the services has teams that do operational red teaming. But, there are other military red teams, like I mentioned before, that try to predict the responses of foreign governments. And so they're doing analytical red teaming, red team gaming, behavioral red teaming, things like that.

Other government agencies and departments have red teams for their particular purposes. Sometimes it's benchmarking. Sometimes they're doing penetration testing. There are research red teams, which we have been one of, who do hypothesis testing and red team gaming and behavioral red teaming, trying to help advance the state of the art of security. And then there seems to be an ever changing number of commercial red teams who perform many of these activities, but seem to concentrate mostly on penetration testing.

Brian Contos: Very interesting. It's interesting, especially doing the analysis of what other organizations or what other governments might think of decisions we make. That sounds like a fascinating area. Taking things back to control systems and SCADA networks and the like, how critical are red teams to ensuring a solid security posture? A lot of us understand the pen testing side for corporate networks and things of that nature, but how critical is this? Is this really an important piece, looking at the cyber end and looking at the actual computing systems behind these types of environments?

Raymond Parks: Well, in the control system world, red teaming has to be carefully controlled, and it's an unusual activity that you have to keep control on. Red teaming for control system environments is not like red teaming an office or an enterprise environment. I rarely, if ever, have been told when I went into an office that if I went to a certain floor of the building that I was going to die if I breathe the air, which has happened to me when I've been red teaming control systems.

Red teams can cause harmful effects in the real world when they're messing with control systems. And not only that, but the real world can cause harmful effects to the red team.

But the real point here is that red teaming can help control system owners determine whether their security investments have been made in the right place and done the right way to prevent a real attack. A lot of people go with compliance standards and regulations. And they're good in their own way. But, standards are one-size-fits-all, and every control system environment has its own unique characteristics.

As part of my work, I'm associated with Sandia's Center for Control System Security. And we teach a concept that we call sustainable security: the best possible security that can be consistently maintained while meeting business objectives with available resources - that's what we defined it as. And red teaming is key to determining the relative value of different security mechanisms and finding what other possible security mechanisms it might be.

We have a process that we developed very recently called Red Team Metrics, which can help determine the best mitigation to stop the most risk of attack. And that way, you can use a red team once you have something established, in order to be able to understand: where should you invest to make yourself as secure as you can be within your resources?

Red Teaming with Ray Parks

Brian Contos: Is cyber security being taken seriously in the industry, in your opinion, and have you seen some successes or failures where it has or it hasn't?

Raymond Parks: Well, like any social or cultural activity, the range of responses to the new wave of cyber security seems to be kind of a normal curve. Some take it very seriously, and some ignore security completely. And somewhere between those outliers, the great majority of industry members and associations have started to work towards securing their control systems.

When I first got involved with control systems, they used to say that control system security was about 10 years behind enterprise system security. Now, I have to say that the gap has narrowed. Control systems are somewhere between five and zero years behind, depending upon their ability to update their systems.

The one critical flaw that control systems still have is that reliance on old systems they can't update without losing control. And so everybody seems to have that one system or that one device that they just can't mess with because they can't find a replacement for it. But, to be fair, I've seen enterprise systems that have the same problem, with some old, vulnerable technology, and they have a lot less excuse for it.

As an aside, I would say that there is a new player that fills the role of being way behind the state of the art. Physical security systems seem to be about 10 years behind the enterprise systems.

You asked about successes and failures. The biggest successes, in my opinion, are all these security standards that various industry groups are developing and adopting. I know a lot of people criticize them as being too lax. And I also know that standards work is just not very flashy or elite in hacker speak, but standards are critical to providing a foundation for security. If you don't have that foundation, there's nothing but inconsistency, and every player is at the mercy of their weakest neighbor. I like to think that standards are the rising tide that floats all boats.

Brian Contos: [laughs]

Raymond Parks: The North American Electric Reliability Corporation Critical Infrastructure Protection suite, NERC CIP, is a particularly good start in this area. I've worked a lot with this, so I'm very familiar with it. And so far as I know - and of course, I'm not an expert on all the standards that are out there - CIP is the only one that addresses the convergence of physical and cyber security.

Also, by focusing the whole process on critical assets that are necessary to the stability of the bulk electric grid, it inherently matches itself to business objectives. That's the other issue, back to my sustainable security point. You've got to be doing things that are going to make your business objectives be met.

Now, as far as individual successes, I have to say we have a limited experience. Sandia doesn't make a business of doing assessments, at least not only assessments. By law, we can't compete if there's a commercial entity capable of doing the work that's out there.

We only do assessments when a customer has done their homework and they've decided that no one else will satisfy their particular need. And even then, we try to limit our control system assessments to situations where we will learn something - do something new, something different, see something we haven't seen before. But, in the process of that, I

Red Teaming with Ray Parks

have seen individual successes in industry. Folks are being proactive, and sometimes they exceed standards in really clever and resourceful ways.

One utility we looked at had implemented 801.1X network access control throughout their control system. And another one used IPsec between all their critical cyber assets, so at least somehow that would help keep some of the external attacks out.

Now, you know, we met during Project Logic, and that's a really great example of a success - a public-private partnership, with the oil and gas industry driving the requirements and the vendors supplying some real, working products, and government and the labs integrating the products to meet those requirements. That was a really good success.

I don't really believe there are any systemic failures within control system security. I know a lot of people, they can say, "Hey, we're just not moving fast enough. We're not getting secure fast enough." But, I don't necessarily agree that that's the case.

I've been involved with computer security, one way or another, for over 30 years. And the initial pace of change in enterprise security was, if anything, slower than control systems. Security is a cultural change, and all cultural change requires time. And if it didn't, it wouldn't be real change. If suddenly, things changed overnight, if everybody said, "Oh, yeah, yeah, we believe in that," then they're probably lying to you.

I said earlier that control systems are catching up. Hey, heck, at this pace, maybe in another five years we'll be asking why enterprise security is so slow to improve.

Brian Contos: Definitely very astute perspectives on this. It certainly has been the mantra, if you will, for many people to criticize the security of control systems. And if that was to surpass the security of standard enterprise, that would definitely shake things up a bit. Ray, this has been fascinating. Time for just some closing thoughts or comments from you?

Raymond Parks: One of the things I would say, in general, about control system securities and red teaming is we need to make sure we keep it up. We've got some momentum going. Things are starting to happen. Cultural change is starting to occur.

And these regulations that are out there, some of which people gripe about not being good enough, and other people gripe about being too hard, they're changing the world view of the people that are involved in control systems and control system security. And I think, over time, everyone will get to the point of getting used to following the rules, and then they'll suddenly realize, "Hey, those rules are here for a reason," and then they're going to start exceeding the rules.

So, I'm hoping that's where we're going to go in the future. And I hope that red teaming is something that people use. There's certainly an increasing interest in using it, for different reasons.

Brian Contos: Well, fantastic stuff, Ray. Thanks so much for joining us on today's podcast.

Raymond Parks: You're welcome. I enjoyed it.

Brian Contos: If you would like to learn more about this subject and Imperva, visit Imperva.com, check out our [blog](#), follow us on [Twitter](#), or send us an e-mail at blog@imperva.com.

Red Teaming with Ray Parks



North America Headquarters
3400 Bridge Parkway
Suite 101
Redwood Shores, CA 94065
Tel: +1-650-345-9000
Fax: +1-650-345-9004

International Headquarters
125 Menachem Begin Street
Tel Aviv 67010
Israel
Tel: +972-3-684-0100
Fax: +972-3-684-0200