All.Net Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

After the Red Team

We used to red team a lot, and we taught a lot of other folks how to red team. I read a lot of disclosures describing and detailing red team activities and presenting results. But what I don't see is anything significant happening after the red team.

A flurry of activity then...

In practice, red teams seem to produce a flurry of activity for a brief period, followed by slipping back into old habits. The situation after the red team is then often worse than it was before, because management now knows some of it's problems. If management doesn't address them there is the potential for far greater liability. So management must act. But what must they do?

The response mismatch

Most information protection leaders in enterprises have lists of things they try to get done and cannot get done because of limited funds or interest by management. When faced with a crisis, real or fabricated, they tell management that the reason they have the problem is that their list of things wasn't done. Then they get emergency funds to do what they wanted to do in the first place.

But what they wanted to do might not address the results of the incident or red team. So there are two approaches:

1. **Control the red team:** By controlling what the red team does and sees, and perhaps by arranging for known weaknesses to be exposed during the effort, the outcomes can be directed down desired paths, resulting in the “discovery” of the known weaknesses.
2. **Control the response:** By controlling the response to the outcomes of the red team, the desired previously unfunded activities get funded, regardless of whether they reflect sensible responses to what was found.

Is it really as illegitimate as that?

To be clear, this can be completely legitimate. A CISO wishing to get funding for known vital changes may engage a red team to demonstrate the problems they know they have. The demonstration allows the protection executive to inform management in a more realistic way than analysis can. Top management often doesn't believe the CISO or views that asserted limitations as unlikely to be exploited. By doing a demonstration, clarity is delivered.

The red team may be completely unaware that they are being driven down a particular path. And many CISOs want to know what the red team discovers, because despite knowing of some weaknesses, they also know they may be missing other ones. Red teams are also often limited in their activities for legitimate reasons. For example, certain sorts of activities may be hazardous, violate regulations, or be outside of the allowable activities of the individual ordering the activity.

Red teaming is a tool used to achieve an end. As such, it is applied to achieve that end.

Is there truly independent and legitimate red teaming?

All red teaming is inherently limited by its nature. Otherwise, it would be no different than actual attack, and likely similarly unproductive and damaging. If it is not intended to achieve a specific goal, it is likely to fail, because there is nothing specific that would meet the goal.

The first step in red teaming is understanding what you want to demonstrate and why. Put another way, red teaming is really a form of protection testing. If you have poorly defined goals, you will likely get undesired results. Before you red team, you should make sure you are getting the right stuff.

How do I get the right stuff?

To get the right stuff, you have to start by understanding what that stuff is. We advise:

- Start by understanding what you are trying to protect and how you are protecting it.
- Devise the test to determine whether what you are doing achieves its intended goal.

Note that this is very different from most current red teaming. Most of the time, what we see is a series of attempted penetrations with some goal (e.g., get into the financial system and prove you can alter a record), some limitations (e.g., using only remote access), and a set of termination criteria (e.g., after you find 3 ways in, stop, and don't spend more than 8 hours).

None of these are likely to be very fruitful for anything but scaring a decision-maker into doing something, likely the wrong thing in terms of having the best protection for the least cost. The reason is that (1) it doesn't answer the question of whether the plan in place is working, and (2) it doesn't help to understand ALL of the ways the mechanisms may be defeated.

- If you cannot figure out what is not protected by your current plan, you should hire someone who knows what they are doing to identify it for you based on the full details of what you are doing instead of hiring someone to blunder about trying things that might or might not work and that could cause real harm.
- If you do a test that doesn't tell you ALL the ways in, it will either find nothing in which case you will think you are safe when you might not be, or it will find some things and you might then notion that you have to fix them or accept the risk, ignoring the things not found.

It's not that analytical approaches are perfect. It's that they typically get a lot more a lot faster and cheaper than testing without clearly defined goals.

Summary

Red teams beget response. The wrong red team or directive begets the wrong response. The right red team with the wrong directive begets the wrong response too. The problem isn't necessarily in the red team, it's often in the directive and/or the response. Proper directives come from understanding what you are testing. Proper response comes from testing the right things and reasonably tying the response to the results. If your goal is to scare management into action, or prove that what you said was vulnerable is, there is also a trust problem that you and your management need to work out.

More generally, the approach here is to start with the objective (i.e., the response) and work your way back to find the right approach to meet that objective.

