The Economics of Cyberwar

Steven M. Bellovin
https://www.cs.columbia.edu/~smb

Abstract

Cyberwar is very much in the news these days. It is tempting to try to understand the economics of such an activity, if only qualitatively. What effort is required? What can such attacks accomplish? What does this say, if anything, about the likelihood of cyberwar?

1 Introduction

Cyberwar is very much in the news these days [2]. At his confirmation hearings as the next NSA director, Admiral Michael Rogers spoke about it extensively [12]. Is cyberwar real, though? Is it a real threat? What sort of effort would it take for a nation to develop a capability?

It is, of course, impossible to answer such questions definitively without access to highly classified information from many different countries. Still, it is possible to approach the question qualitatively and get some rough ideas.

It is important to realize that while cyberweapons are based on many of the same principles as ordinary criminal malware, they can vary very much in detail. A modern bomber differs greatly from a commercial airliner, even though both rely on wings, jet engines, and the laws of physics. The same is true in cyberspace. While we do not know for certain what cyberweapons look like, the few advanced attack tools that have been seen in the wild—Stuxnet [3], Gauss [5], Flame [4], and more—give us some idea what can be done by highly skilled, well-funded professionals.

It is also important to ask how effective such weapons can be, and under what conditions. Are they a serious threat? Are nations at risk of a catastrophic first strike? Again, there are no certain answers, but we can draw some tentative conclusions.

2 Effort

At one end of the scale, it is clear that developing "weaponized malware" is not a task for amateurs. Stuxnet, the first known piece of malware that was clearly developed by a nation-state [13], is an interesting case study. It used four different "0-day" attacks, attacks that by definition were not known to the vendor. That alone shows the attacker's resources: 0-days, though not exceptionally rare, are still relatively uncommon. For one party to have four of them—four that it was willing to use in a single attack—suggests that there were many more in its arsenal. There is another conclusion that we can draw: it took that many 0-days to attack such a hard target. If that observation is more generally true, that it takes a lot of effort to launch a successful cyberattack against a well-defended target—in this case, a uranium centrifuge plant whose computers were isolated by an "air gap"—it suggests that cyberwar is not as easy as some would portray it, and that World War C cannot be launched by a couple of bored teenagers who have run out of pizza.

Flame [4] showed sophistication of a different type: it used a previously-unknown cryptanalytic technique as part of its attack. Again, this is something that is beyond the ability of most governments, let alone of terrorist groups.

To be sure, more ordinary techniques will often suffice. Press reports based on the Snowden documents suggest that the NSA has hijacked ordinary botnets and used them to gather intelligence [10]. Botnets rarely affect secure networks, except by accident; consequently, a militarized bonnet is more useful for lesser targets. Still, there are plenty of those. Much of the cyberespionage attributed to China appears to have been carried out with tools of this grade. Indeed, the attack on the New York Times (itself generally blamed on the Chinese [9]) was apparently carried out by a "B-team" of attackers; their poor operational security practices suggest a less-than-professional effort. It should be noted even military targets can be penetrated this way, especially if their networks aren't properly run. A recent penetration of a U.S. Navy network, attributed to Iran [6], was "enabled by vulnerabilities that were discovered and exploited in older systems and network architecture"; there was apparently little internal security, and a common attack technique known as "SQL injection" was apparently used. These attackers may or may not have been extremely skilled, but the initial penetration was of a type very commonly used by ordinary criminals.

There is thus a wide range of abilities needed; a lot depends on the target of the attack. Flipping that around, defense can succeed, at least against all but the strongest attackers. This is in marked contrast not only to the threat of nuclear war, where no defenses seem feasible, but also to the common wisdom about cyberwar.

3 Effectiveness

Assume that an attack can be launched. How damaging might it be?

Libecki has published one of the more thorough analyses [8]. He suggests that such software is better seen as a tactical weapon rather than as a strategic one. He bases this largely on the unpredictability of the effect of an attack; a minor change in configuration, software version, patch level, etc., of the target network can render the attack futile.

There is another factor to consider: the response of the defenders. If a system or network crashes, the operators and system administrators will try to bring it back up. They'll use backup media, monitors to detect ongoing attacks, changes in firewall rules, etc. These efforts won't always succeed, of course, but they won't always fail, and the more time the defenders have, the more likely they are to prevail. Again, this demonstrates the role of defense. It also underscores Libicki's thesis that cyberweapons are better relegated to a tactical role.

There is one important exception, though: if a computer is connected to a physical system—SCADA systems are the obvious example—serious physical damage can be done before the attack is repelled. That is, the target network may be back up after a few hours or days; in the interim, however, the chemical plant has exploded or the nuclear reactor has melted down. A consequence of this is that such systems require far more protection than ones where the effects of misbehavior can be corrected. The financial system, as critical as it is, would seem to fall into the category of systems where rapid recovery is an acceptable second choice to a strong defense.

All that said, a sophisticated attacker with ample resources may be able to transcend the tactical role. If intelligence is good enough to know about configuration changes, the attacker's methods could change in response. (I note that excellent intelligence is necessary to attack any well-defended target. Indeed, that is one of the most striking aspects of Stuxnet and Gauss [5]: they were very precisely aimed, which in turn demonstrated how much was known about their target networks. Regardless of how this information was gathered—spies, cyberspies, or good analysis—the attackers had the information and used it.) Similarly, recovery efforts can be thwarted by adaptive attacks, or by attack software that wormed its way into the target systems so long ago that all useful backups are themselves infected.

This approach—ongoing, long-term involvement with possible targets—is expensive. It's especially expensive if done for a significant number of targets; this of course means that a poorer attacker is less able to do it. Conversely, a nation with the resources and skills to automate the process, to continuously gather detailed intelligence, and to "prepare the battlefield" will have a significant advantage.

Long-term involvement poses its own set of risks for the attacker, because it permits different defensive strategies. Not unreasonably, today's intrusion detection systems are designed to detect today's attacks, whether they're direct or "low and slow" [14]. Generally, though, these systems are intended to react quickly, as soon as there is hard evidence of a problem.

A penetration designed for long-term monitoring does not have to be found immediately; rather, it has to be found before it's used. This means that expensive scans can be employed to find the problem, because they can be done infrequently. Alternatively, routine reinstallation of the operating system and applications, across the board, can be done at irregular intervals, to try to flush out the problem. Long-term monitoring of communications, especially after exactly the sort of configuration changes that would be of interest to attackers, might also prove profitable.

4 Conclusion

Cyberwar is not an all-or-nothing proposition. The biggest single variable is probably the total resources of the attacker, with its technical sophistication close behind. A lesser attacker than a superpower—a smaller country, a terrorist group, a gang of idle 20-somethings—can undoubtedly cause damage, even serious damage; however, the scale is not likely to be catastrophic, and the actual effects will be hard to predict in

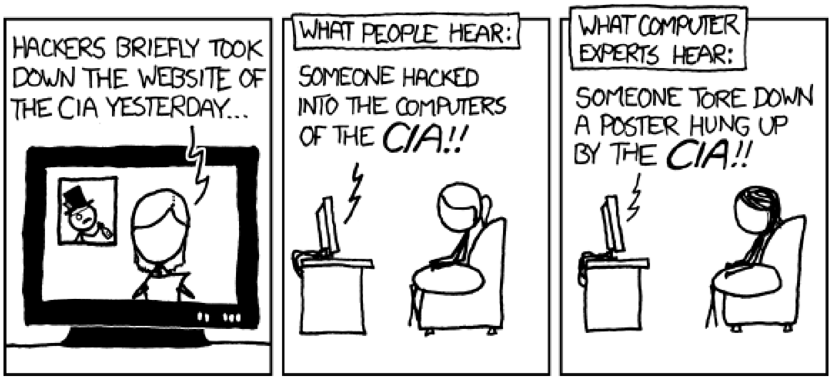
HACKERS BRIEFLY TOOK DOWN THE WEBSITE OF THE CIA YESTERDAY...
WHAT PEOPLE HEAR:
SOMEONE HACKED INTO THE COMPUTERS OF THE CIA!!
WHAT COMPUTER EXPERTS HEAR:
SOMEONE TORE DOWN A POSTER HUNG UP BY THE CIA!!

Figure 1: The significance of denial of service attacks. (`https://xkcd.com/932/`)

advance. By the same token, strong defenses, in the form of system configurations, good system administration, proper backups, and accurate intrusion detection systems [7], are a good counter. It remains unclear who would win if a high-end offense were to take on a high-end defense.

As noted earlier, this suggests a different policy approach than many countries seem to be following. Rather than stressing offense, countries should concentrate on defense and intrusion detection. It's more scalable (one well-written web server can run on millions of computers); additionally, a good defense will likely prevail against almost all attackers.

Some would say that distributed denial of service attacks (DDoS) are an exception. Indeed, such attacks have been carried out as part of a conflict, e.g., in Estonia [11]. Although these attacks can be serious, they do not cause physical damage to SCADA-controlled chemical plants. More often, they're similar in effect to vandalism (see Figure 1). Furthermore, improving our defenses—that is, protecting *all* computer systems, world-wide—will deny the would be DDoSer the wherewithal to carry out the attack. Helping everyone helps ourselves.

Finally, it should be noted that wars do not break out simply because one side has a new weapon [1]. Rather, they occur because of underlying enmity or tension. A "digital Peal Harbor" is likely to occur under the same circumstances as the kinetic Pearl Harbor attack of 1941: when two nations are already at the brink of war.

References

[1] Steven M Bellovin. "Military Cybersomethings". In: *IEEE Security & Privacy* 11.3 (May 2013), p. 88. URL: `https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6521321`.

[2] Richard A. Clarke and Robert K. Knake. *Cyber War: The Next Threat to National Security and What to Do About It*. New York: Ecco, 2010.

[3] Nicolas Falliere, Liam O Murchu, and Eric Chien. *W32.Stuxnet Dossier*. Symantec Security Response. Version 1.4. Feb. 2011. URL: `http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf`.

[4] Dan Goodin. "Crypto Breakthrough Shows Flame Was Designed by World-Class Scientists". In: *Ars Technica* (June 2012). URL: `http://arstechnica.com/security/2012/06/flame-crypto-breakthrough/`.

[5] Dan Goodin. "Puzzle box: The quest to crack the world's most mysterious malware warhead". In: *Ars Technica* (Mar. 14, 2013). URL: `http://arstechnica.com/security/2013/03/the-worlds-most-mysterious-potentially-destructive-malware-is-not-stuxnet/`.

[6] Siobhan Gorman. "Navy Hacking Blamed on Iran Tied to H-P Contract". In: *Wall Street Journal* (Mar. 6, 2014). URL: `http://online.wsj.com/news/articles/SB10001424052702304732804579423611224344876`.

[7] Wenke Lee and Salvatore J. Stolfo. "Data Mining Approaches for Intrusion Detection". In: *7th USENIX Security Symposium*. San Antonio, Texas, 1998. URL: `http://static.usenix.org/publications/library/proceedings/sec98/lee.html`.

[8] Martin C. Libicki. *Cyberdeterrence and Cyberwar*. Monograph MG-877. Rand Corporation, 2009. URL: `http://www.rand.org/pubs/monographs/MG877.html`.

[9] Mandiant. *APT1: Exposing One of China's Cyber Espionage Units*. White paper. 2013. URL: `http://www.mandiant.com/apt1`.

[10] Joseph Menn. "NSA 'hijacked' criminal botnets to install spyware". In: *Reuters* (Mar. 12, 2014). URL: `http://www.reuters.com/article/2014/03/12/us-usa-security-nsa-botnets-idUSBREA2B21420140312`.

[11] Jason Richards. "Denial-of-Service: The Estonian Cyberwar and Its Implications for U.S. National Security". In: *International Affairs Review* (Apr. 2009). URL: `http://www.iar-gwu.org/node/65`.

[12] David E. Sanger. "N.S.A. Nominee Promotes Cyberwar Units". In: *New York Times* (Mar. 11, 2014). URL: `http://www.nytimes.com/2014/03/12/world/europe/nsa-nominee-reports-cyberattacks-on-ukraine-government.html`.

[13] David E. Sanger. "Obama Order Sped Up Wave of Cyberattacks Against Iran". In: *New York Times* (June 2012). URL: `http://www.nytimes.com/2012/06/01/world/middleeast/obama-ordered-wave-of-cyberattacks-against-iran.html`.

[14] Salvatore J Stolfo. "Worm and attack early warning: piercing stealthy reconnaissance". In: *Security & Privacy, IEEE* 2.3 (2004), pp. 73–75.