Conference on Cyber Conflict
Proceedings 2010
C. Czosseck and K. Podins (Eds.)
CCD COE Publications, 2010, Tallinn, Estonia.

PINPRICK ATTACKS, A LESSER INCLUDED CASE?

Antoine LEMAY[a,1], José M. Fernandeza, Scott Knight[b]

[a]École Polytechnique de Montréal, Montreal, Canada,
[b]Royal Military College, Kingston, Canada

Abstract: Defense has always been tailored to threats; this prevents wasteful resource spending and strategic surprise. However, with the introduction of asymmetric warfare techniques, including cyberwarfare, defending against all threats has become impossible. To deal with this problem, the notion of "warfare spectrum" was introduced. At one end of the spectrum stands complete peace, at the other end high-intensity kinetic warfare. The theory behind this was that a force trained for high-intensity would be able to deal correctly with "lesser included cases" in the spectrum. This way of thinking has also been applied to cyberwarfare and critical infrastructure defense.

In the literature, we can notice a definitive focus on preventing a "cyber Pearl Harbor" or "cyber 9/11", meaning an unforeseen, devastating attack. Alternatively, following the events in Georgia and Estonia, the protection from massive coordinated denial of service was also considered. Still, both of these scenarios sit in the high-intensity spectrum of cyberwarfare. However, in our analysis, we have found that low-intensity cyberwarfare could be as devastating and cannot be considered a "lesser included case" of high-intensity cyberwarfare, contrary to what the "warfare spectrum" theory dictates. In particular, we present the "pinprick attack" scenario, where the goal of an attacker is to produce long-term damage by the accumulation of large numbers of low damage attacks.

In our paper, we demonstrate why new solutions are needed to defend against our scenario. First, we illustrate how a "Clausewitzian" definition of warfare limits the kind of responses that are available to the target of such an attack.

[1] École Polytechnique de Montréal, 2500 chemin de Polytechnique, Montréal, H3T 1J4, Canada, Email: antoine.lemay@polymtl.ca.

Because no formal declaration of war is made, responsibility for defense will rest on the private sector and not military institutions. Then, we see that existing defense solutions, such as data aggregation by national agencies and generalized vulnerability reduction, fare poorly against the pinprick attack scenario because the damage threshold is kept small and because the attack's breadth is very large. Finally, we present some ideas to counter "pinprick attacks". Notably, we mention the optimization of defensive solutions to cover a wider range of threats (our own research project) and regulatory economics (field for future work).

Keywords: cyber warfare, asymmetric warfare, critical infrastructure protection

INTRODUCTION

Defense has always been tailored on threats; this prevents wasteful resource spending and strategic surprise. However, with the introduction of asymmetric warfare techniques, including cyberwarfare, defending against all threats has become impossible and defenses are focused on likely threats. This leaves holes that can be exploited by new attack forms. In particular, we will analyze how the assumption that a cyberwarfare opponent would use a high-intensity form of cyberwarfare creates holes where a low-intensity form of cyberwarfare can thrive.

We will start by looking at how governments plan to solve the cyber security problem of critical infrastructure. This will allow us to extrapolate the attack scenarios that are considered high threat. We will then compare this scenario with current military thinking in order to confirm our extrapolation. We also analyze the limitations inherent to the scenario. Finally, we present an attack form, the pinprick attack, which uses these limitations to maximize the damage it can cause and we offer avenues for future research that would enable defenders to defeat our attack.

1. CRITICAL INFRASTRUCTURE PROTECTION

A lot of effort has been invested to bolster cyber security. A group of experts mandated by the Center for Strategic and International Studies (CSIS) argued again in 2008 in their Cyber Security for the 44th Presidency report that "cyber security is now a major national security problem for the United States" (CSIS, 2008). Within that major national security problem lays the problem of securing critical infrastructure against attack. Various solutions have been proposed to reduce the risk associated with cyber attacks on the critical infrastructure. However, these solutions are based on unconscious strategic assumptions that might prove not to be true.

In this section, we will look at the two most common propositions to reduce cyber security risks on the critical infrastructure. We then analyze the solutions to extract the scenarios they are most useful against. Based on this analysis, we draw conclusions about the strategic assumptions that drive the efforts to reduce risk.

1.1 VULNERABILITY REDUCTION

The most common solution to reduce the risk for the critical infrastructure is some sort of vulnerability reduction program. The 2003 National Strategy to Secure Cyberspace has two national priorities addressing this issue. Priority II (a national cyberspace threat and vulnerability reduction program) addresses technical vulnerabilities while Priority III (a national cyberspace security awareness and training program) addresses human vulnerabilities (Department of Homeland Security [DHS], 2005). Various methods have been employed to attain these goals. One example of vulnerability reduction program is the North American Electric Reliability Commission (NERC) Critical Infrastructure Protection (CIP) standards (NERC, 2010) that are required to be met by January 2010. The idea behind this strategy is that once vulnerability has been reduced, an opponent will not have any opportunity to attack.

The underlying assumption behind the concept of generalized vulnerability reduction is that it is possible to reduce your vulnerability enough to make attacking you inefficient. It is clearly not possible to reduce the vulnerability over the entire attack surface. As Welander shows in his review of cyber security for the industrial control sector (Welander, 2009), skilled and motivated attackers, such as spies and extortionists, tend to use more sophisticated attack strategies. In particular, highly committed opponents can afford to use a strategy of systematic probing for vulnerabilities. In fact, they can also attempt to induce vulnerability in the target by finding undisclosed vulnerabilities or by distributing Trojan horses or backdoors for example. As skill and motivation increase, it becomes increasingly costly to reduce vulnerability to a point where no risk exists. In that light, the implied objective of national vulnerability reduction programs is to address the lower left quadrant of Figure 1, i.e. widely known vulnerabilities affecting your industry in general. This is even truer if the private sector is to assume the costs of vulnerability reduction as in the case in NERC CIP standards. Because the private sector is profit-driven and has no vested interest in national security, market forces will drive the private sector to minimal compliance. That minimal compliance will be aimed at defeating casual attackers, which is possible to do at reasonable costs, and not highly trained and motivated attackers, which are an unlikely threat and very costly to defend against.

1.2 DATA CORRELATION

The other solution that is most often proposed is the creation of a national agency to collect and correlate data. This can take various forms. For example, in the Department of Homeland Security (DHS) report on the National Strategy to Secure Cyberspace, Priority I is a “National Cyberspace Security Response System” (DHS, 2003). In the report for the 44th Presidency, the authors ask that the president “reinvent the public-private partnership” (CSIS, 2008). This is usually done by the creation of Computer Emergency Response Teams, or CERTs as described in the DHS press release detailing its activities in regards to the National Strategy to Secure Cyberspace (DHS, 2005). Once established, the CERTs share information with the various government agencies and the private sector. The idea being that the global situational awareness obtained through the centralization of information and the established relations with various actors will allow the CERT to successfully coordinate efforts to diffuse a crisis. This model is widespread even if only US sources are presented. We can find CERTs in the US, in Canada, Australia, Estonia and even in non-NATO countries such as Russia.

If one assumes that the CERT model works as designed (and the various improvements suggested in the report to the 44th Presidency suggest that it may still require improvement), the CERT model itself is based on a critical assumption. It is assumed that centralization of data will produce an increased situational awareness that can be turned into a defensive advantage. The only scenario where that assumption is likely to prove correct is in the case of a concerted effort by an attacker to target a variety of CERT partners. For example, an opponent coordinating DoS attacks on government servers, banks and television networks would be able to be easily correlated by a CERT and actions could be taken to deal with the situation as a whole instead of in isolation. However, in order to make such a correlation, it is necessary to have some sort of link between the attacks such as a temporal link (e.g. after a political event). Other types of linkage are possible, but may not enable a CERT to produce a coordinated defense. For example, a series of attacks using the same methodology over a long period of time could be eventually correlated, but it would likely be too late for a response.

1.3 STRATEGIC ASSUMPTIONS

As we have seen, proposed solutions to reduce the risk to critical infrastructure are based on specific risk scenarios. In the case of vulnerability reduction, we want to reduce the exploitation of low hanging fruits vulnerabilities by unskilled attackers. In the case of centralized data correlation, we hope to be able to detect and respond to correlated attacks. This kind of attack footprint can be associated with a limited number of strategic scenarios.

The first scenario is the asymmetric opponent. In this scenario, an opponent decides to target your infrastructure with a massive cyber attack to make you hurt as much as possible. This can be used as a support for deterrence much in the same way as other asymmetric warfare tactics (e.g. insurrection) are attempted. The attacks in this scenario are performed by an inferior opponent. They are likely to be limited in terms of skill because of the limited resources that can be deployed by the inferior opponent who may not possess highly trained assets that can exploit less widely known vulnerabilities or does not have a large amount of time to induce vulnerabilities or perform exhaustive vulnerability searches. Also, the attacks are likely to be correlated in time (linked with specific deterrence event) and space (originating from the same region). Similarly, coordinated effort is likely to be worthwhile because of the high correlation.

The second scenario is the use of cyberwarfare to support military operations. The most common example is the use of cyberwarfare to perform command and control warfare. In that example, the cyber attacks are heavily correlated in time (with conventional warfare operations) and targeting (command and control assets). Response can also be easily centrally coordinated as part of a military response. Because it is linked with military operations, a high tempo can be expected. In that sense, limited use of exhaustive vulnerability research and research for new vulnerabilities is not likely to happen once operations start. In that sense, the attack footprint would be similar to the asymmetric opponent scenario.

In both of these cases, we are dealing with a clear opponent and a high tempo of cyber attacks. As such, both of the scenarios can be considered high-intensity cyberwarfare. But are there low-intensity cyberwarfare scenarios?

2. INTENSITY IN CYBERWARFARE

Based on the solutions that are proposed to reduce the risk for critical infrastructure, one might extrapolate that our main concern is high-intensity warfare scenarios. In this section, we see how this fits conventional western military thinking and the limitations of this view.

2.1 WARFARE SPECTRUM

Western military doctrine is significantly influenced by the works of Von Clausewitz. In particular, that war is the continuation of politics by other means. This led to the development of the "spectrum of warfare", described in various doctrine documents such Canada's Army (National Defence Canada, 1998) and Land Operations (National Defence Canada, 1998). Figure 1 illustrates the concept.

As we can see, operational military means are only employed in times of conflict or war. In that mindset, it is normal that cyberwarfare would be employed in the same conditions. These conditions dictate how force is used, even for cyberwarfare. In a condition of war, the goal is usually to bring a quick end to the conflict. As such, there is no incentive to limit the damage you are doing to the enemy. This is consistent with the attack profiles for high-intensity cyberwarfare presented earlier.

Because of the dangerous nature of the warfighting end of the spectrum, modern armed forces are trained first and foremost to deal with combat operations. The rationale is that if you are trained for the difficult, you will excel at easier tasks.

<table>
<tr><th>Peace</th><th colspan="2">Conflict</th><th>War</th></tr>
<tr><td colspan="3"><b>Military operations other than war</b></td><td></td></tr>
<tr><td colspan="3">Strategic military response</td><td><b>Warfighting</b></td></tr>
<tr><td colspan="3"><b>Non-combat operations</b></td><td></td></tr>
<tr><td colspan="2">Operational military means</td><td colspan="2"><b>Combat operations</b></td></tr>
</table>

Figure 1. Spectrum of warfare

This is confirmed by Canadian doctrine. In the Land Operations publication (National Defence Canada, 1998) we read that "combat capable forces are flexible enough to adapt to the requirements of **non-combat operations**" (original emphasis). In other words, non-combat operations are lesser included cases of combat operations. By following this thinking in cyberwarfare, it makes sense to concentrate on defending for high-intensity cyberwarfare.

2.2 LIMITATIONS

The main limitation of the traditional western military thinking is that military response is not triggered until the conflict has been escalated. Typically, some sort of declaration of war or act of war is required. In the cyberwarfare world, this would require a successful correlation of the attacks before committing to an organized response. If the correlation cannot be made, the defense framework that is in place (e.g. CERT teams, government agencies, etc.) cannot be used. Also, because time is one of the primary factors that drive attack correlation, low-intensity warfare is unlikely to be successfully correlated as “warfare”. This is not a problem when dealing with other nations that are following the same set of principles for warfare and politics, but can become a problem when dealing with countries (or organizations) that do not.

The intense competition between classical Chinese states as illustrated in Chinese military classics (Sun Tzu (2006) and Sawyer (1993)) offers a great example of a diverging theory for what constitutes warfare. Everything your state gains at the expense of other is ultimately a strategic advantage that you will be able to use later and thus is, in essence, warfare. This way of thinking is still present in modern Chinese military literature. For example, in the book “Unrestricted Warfare”, Liang and Xianshui (1999) argue that multiple forms of warfare such as financial warfare, trade warfare and cyberwarfare could play a major role in wars of the future. Obviously, the role of these alternate forms of warfare is to diminish the fighting strength of a nation by attacking the national assets that support the military establishment. Naturally, no nation would allow itself to be attacked in that fashion.

This leads to another limit on the concept of high-intensity cyberwarfare. There are inherent limits to the damage you can cause to any opponent that has the means to defend itself. The first limit is the ability for the target to “pull the plug” or disconnect his network from yours. Even the Internet requires a backbone, which can be deliberately partitioned by cutting a limited number of points (for example the endpoints of oceanic cables (Internet’s Undersea World, 2010)). So, if you are facing a rational opponent, the damage he can inflict on himself by pulling the plug (and whatever you can sneak in before he does) is the upper bounds to the damage you can inflict. If he assesses that you can do more damage to him than the damage of pulling the plug, he will disconnect, and if you can’t he will accept your damage. The second limit is the ability for the target to escalate. To illustrate, let us consider what would be the US response to an enemy trying to disable a vital strategic asset such as the US nuclear command and control system. We can easily extrapolate that this would provoke a significant response using a broad spectrum of means.

By taking a low-intensity approach, it is possible to abuse these limitations to create a new cyberwarfare threat.

3. PINPRICK ATTACKS

Pinprick attacks are an illustration of what can be done with low-intensity cyberwarfare. With Pinprick attacks, the trick is for the attacker to lead the defender into believing he is facing unconnected single instances of small attacks. This is done by staying under his correlation threshold. It is similar to the practice of "slow slicing" or "death by a thousand cuts" in the sense that you do not perform a single crippling attack, but instead a collection on non-crippling attacks whose effects add up to create the crippling effect.

3.1 DESCRIPTION

In our pinprick attack scenario, individual damage per incident is low. It is therefore ill suited to attack hardened targets built with resilience in mind such as military communications. However, because it is a long-haul strategy, we can perform attacks on select points which will yield good results. The specific targeting of ball bearing factories by US bombers in World War II is an example of operations designed to destroy a fighting capability without actually directly targeting military hardware. Can such an operation be carried out in a cyberwarfare context? RAND's publication "Measuring National Power in the Postindustrial Age" (Tellis et al, 2000) offers us some insight into how this could be done. This report presents a methodology to evaluate a nation's power using more than military power as the sole criterion. In the RAND model, combat proficiency is a result of the combination of strategic resources and the capability to convert these resources into military power. The easiest example is the case of military technology. A country with rich resources into terms of knowledge and money (strategic resource) can transform this resource in military technology through its military-industrial complex (conversion capability). Because we are talking about a combination, affecting either the resources or the conversion capability will result in a decrease in military power. We could present our "death by a thousand cuts" scenario as gradually injecting grains of sand into a complex clockwork mechanism in order to make it stop, or at the very least run less efficiently.

Defense from this scenario, in western countries, is mostly under the control of the private sector. For example, privately owned banks control most of the financial system, privately owned power companies supply the power, privately owned companies produce most of the technology and hardware used by the military. The goal of these companies is to make profit. This objective is usually incompatible with spending money to defend against an unlikely scenario (e.g. cyberwarfare). Increased spending for cyber security can even be detrimental to the health of a company. After all, if your costs are higher than those of your competition because of high security measures, customers will buy your competitor's products. This breeds a vulnerability-rich environment that drives the costs of creating an attack operation down even in the face of government-mandated vulnerability reduction programs. Attackers have all the time they need to perform exhaustive searches for vulnerabilities because the attack follows a deliberately slow tempo. This gives a determined attacker the agility required to attack only targets of opportunities and to follow the path of least resistance and pick the low hanging fruits. In that sense, a vulnerability reduction program does not offer adequate protection against pinprick attacks.

An important aspect of pinprick attacks is to keep the defender unaware that the attacks he is seeing are part of a coordinated strategy. As long as he is not able to correlate the attacks, there is no theoretical limit to the amount of damage you can inflict. This can be explained by the fact that, compared with each incident in isolation, the cost of coordinated response will always be higher than the incident's damage. For example, if you find a Trojan horse on a military contractor's computer, you clean it and try to assess the damage. If you find one on someone else's computer the next week, you will do the same. However, if you find a Trojan on the computers of all the military contractors, you might take more active measures to stop whatever is going on. So, by design, pinprick attacks are difficult to defend against by centralized data correlation agencies such as CERTs.

3.2 EXAMPLE

Because pinprick attacks reside in the low-intensity part of the spectrum, they are not well suited for what we consider warfare scenarios, which require speedy conflict resolution. However, it is ideally suited for competition between near peers where one of the peers wants to slow down the progress of his other peers to catch up with them or increase its advantage.

Let us consider the fictional scenario where the countries of Alpha and Beta are near peers. However, the people of Alpha possess a significant advantage in technology over Beta. This advantage in technology allows the military of Alpha to hold a strategic advantage over Beta's military force, even if both are similar in other aspects. If Beta were to pursue a high-intensity cyberwarfare strategy, Alpha could respond by pulling the plug and escalating to a military conflict where Alpha has the advantage. This course of events is therefore detrimental to Beta. However, Beta can instead decide to be patient and use pinprick attacks, slowly but methodically launching attacks to undermine the confidentiality around Alpha's technology. Beta can sum up the benefits of all his attacks (plans captured by a Trojan, information recovered from a stolen USB key, communications intercepted on the wire, etc.) to catch up with Alpha in technology and negate Alpha's strategic advantage. It is unlikely that Alpha would recognize that the various incidents are connected to a coordinated effort by Beta to negate a military advantage because individual incidents only cause limited damage.

3.3 COUNTERING PINPRICK ATTACKS

As we have seen previously, the solutions that are currently proposed to deal with cyber threats are not really appropriate to deal with pinprick attacks. In order to defend effectively against them, new solutions are required.

The ideal solution would be to possess the means to correctly correlate attacks, but this is very difficult. After all, the attacker can set the tempo to whatever value allows him to evade detection (although there is admittedly a value under which the tempo would be too low to produce significant damage). We must therefore concentrate on vulnerability reduction. Again, as we have seen, this can also be a daunting task. However, unlike correlation, defenders have the levers of technology and economics to tackle the problem. In both cases, the goal is not to completely reduce the vulnerability, but instead to reduce the damage to the investment ratio of the attacker.

This can be achieved by having better technology. If, with the same market constraints, we can provide better security, we will blunt the attacker's advantage. If we manage to build security devices that are cheaper, implementing adequate security will prove less of a burden on the private sector. It will then be possible to ask more security of the private sector. Similarly, finding ways to optimize the efficiency of existing technology is another avenue that can be pursued. In particular, finding ways to use existing technology to extend the threat coverage could prove to be an interesting field of research in that regard.

The other option to increase the overall security is to change the market constraints. A tool governments have at their disposal is regulatory economics, e.g. by providing subsidies to critical infrastructure operators to upgrade their security. Another example would be the creation of penalties if some level of security is not achieved as is the case in the NERC CIP standards (NERC, 2010). While our research group is not focused on economics, this field could prove to be fruitful for further research.

4. CONCLUSION

In this paper, we have analyzed the solutions that are more commonly proposed to deal with the cyber security of the critical infrastructure. In particular, we have seen that national programs of vulnerability reduction are mostly successful in reducing the vulnerabilities used by unskilled attackers. As for centralized correlation of data, we have seen that it requires distinguishable patterns in the attacks to be successfully correlated. More importantly, we argued that successful correlation is required for a coordinated defense. These limitations reveal the underlying assumption that the expected opponent will use some form of high-intensity cyberwarfare. While that assumption is reasonable for an opponent following a military doctrine based on von Clausewitz's writings, we cannot assume that all opponents would adhere to such a philosophy.

To prove that low-intensity cyberwarfare is possible, we have proposed the "pinprick attack" scenario where an opponent launches a series of attacks too small and too distant to be successfully correlated. Because the attacks cannot be correlated, a nation cannot offer a coordinated response such as escalating the conflict to a field more advantageous for the defender, such as conventional warfare, or such as "unplugging" from the network. The attacker can then endlessly repeat his attacks to cause a "death by a thousand cuts".

Because current defensive strategies are not well adapted to deal with pinprick attacks, future work is required to bolster defenses. In particular, research to reduce the financial burden of security for critical infrastructure operators is an avenue that our research group pursues. Another promising avenue of research would be the use of regulatory economics to change the market forces that drive the critical infrastructure operators to the lowest common denominator.

REFERENCES

- Author unknown, 2010. *The Internet's Undersea World*. Internet: http://image.guardian.co.uk/sys-images/Technology/Pix/pictures/2008/02/01/SeaCableHi.jpg, [Feb. 2010]
- CSIS Commission on Cyber Security for the 44th Presidency, 2008. *Securing Cyberspace for the 44th presidency*. Internet: http://csis.org/files/media/csis/pubs/081208_securingcyberspace_44.pdf, December 2008 [Feb. 2010]
- Department of Homeland Security, 2003. *The National Strategy to Secure Cyberspace*.
- Department of Homeland Security, 2005. *Fact Sheet: Protecting America's Critical Infrastructure – Cyber Security*. Internet : http://www.dhs.gov/xnews/releases/press_release_0620.shtm, Feb. 15 2005 [Feb. 2010]
- Liang, Q. and Xianshui, W., 1999. *Unrestricted Warfare*. PLA Literature and Arts Publishing House
- National Defence Canada, 1998. *Canada's Army*.
- National Defence Canada, 1998. *Conduct of Land Operations*.
- North American Electric Reliability Commission, 2010. *Critical Infrastructure Protection standards*. Internet : http://www.nerc.com/page.php?cid=2|20, [Feb. 2010]
- *T'ai Kung's Six Secret Teachings* in Sawyer, R. D., 1993. *The Seven Military Classics of Ancient China*. Westview Press, USA.
- Sun Tzu, 2006. *The Art of War*. translated by Griffith, S. B., Blue Heron Books, Canada
- Tellis, A. J. et al, 2000. *Measuring National Power in the Postindustrial Age*. RAND, USA.
- Welander, P., 2009. *Cyber Security*. Control Engineering, vol. 56, no 1, January 2009, p. 41.