SAND2007-0550J

Identifying and Defeating Blended Cyber-Physical Security Threats

G. Wyss, P. Sholander, J. Darby, and J. Phelan
Sandia National Laboratories,
PO Box 5800, Albuquerque, NM, 87185-0757, gdwyss@sandia.gov

Introduction

Critical infrastructure security professionals have long recognized that their systems are potentially vulnerable to both physical and cyber attacks. Infrastructure owners have invested large sums to identify and remediate potential vulnerabilities in physical protection systems and computer network operations. Yet, in many ways, physical security and cyber security remain separate and independent disciplines. Their analyses are performed by separate teams and documented in separate reports, with remediation plans generated and implemented by separate organizations with very different cultures (e.g., “geeks” versus “cops”). However, modern infrastructure systems consist of both physical and cyber components that interact with one another in complex ways. The same is true of modern security systems, which are composed of physical barriers and detectors, many of which are under computer control. Interactions between physical and cyber security are recognized in the popular media (e.g., the movie “Ocean’s 11”), but those interactions are often minimized in our own physical and cyber security analyses. This paper describes a straightforward way to evaluate “blended” attack pathways where the adversary uses both physical and cyber attack tactics.

Types of Attacks

Each infrastructure attack is categorized traditionally as either a physical attack or a cyber attack. In a physical attack, an attacker gains physical access to a physical asset in the infrastructure system in order to damage it, disable it, steal it, or use it in an undesirable way. The adversary uses force, stealth, or deception to disable or bypass access controls, and completes the attack either by manipulating the system (e.g., turning valves, opening circuit breakers) or by causing physical damage to its components. Protection against physical attacks focuses on the timely detection, interruption and neutralization of adversary activities. Security system designers want to detect adversaries reliably and in time to mount an effective response that first interrupts and ultimately neutralizes the adversary. To buy time for responders, security system designs include elements that delay the adversary’s progress and thereby extend the timeline for access and attack. The physical security of a facility can be evaluated using simple tools that consider industry-recognized best practices or with more advanced tools that construct scenario-specific adversary and defender timelines to estimate the likelihood of a timely security response [1, 2]. The goal is to produce a security system that has balanced physical protection for all potential adversary scenarios. To use an analogy, we want to avoid a system that has very strong doors but wide-open windows.

In a cyber attack, adversaries accomplish their goal by cyber manipulation of the system (e.g., turning components off, changing set points, disrupting flows, deleting data, etc.) without ever gaining physical access to the affected component. For commodity delivery infrastructures, cyber attacks are usually of concern only if undesired consequences can be caused by cyber manipulation of a physical asset. For other infrastructures (e.g., banking), non-physical consequences such as information modification or deletion are also of concern. Protection against cyber attacks focuses on ensuring that unauthorized users cannot access the system, that authorized users' capabilities to cause damage are limited, and that system restoration can be accomplished quickly. These goals are accomplished using features such as user authentication, access control, encryption, monitoring, integrity checking, redundancy, and disaster recovery planning (including system backups). Many tools are available to help analysts implement industry-recognized cyber security best practices [3], and network analysis and attack graph methods are available for more in-depth assessments. [4]

An important aspect of both physical and cyber security system design is to understand the types of adversaries and attacks against which the system will defend. Clearly one cannot build infinitely strong defenses in either the physical or cyber security realms. Therefore, the security analysis team must work with management, first, to understand the types of adversaries that may wish to attack your system, and second, to fashion appropriate responses should such an attack occur. Appropriate responses may include security to thwart the attack, redundancy or recovery measures to minimize the impact of the attack on the system, mitigation measures to minimize the public consequences of system failure, and even insurance to minimize the impact of the attack on the company. For many infrastructures, detection and mitigation of adversary actions may be more cost effective than increased security.

The description of the adversary groups that the security system must defend against is sometimes called a Design Basis Threat or Plausible Threat Envelope, and should include factors such as the number of adversaries, areas of expertise, types of equipment, financial resources, and even tactical sophistication and level of motivation for the attacking force (e.g., are they willing to die? Willing to be exposed? Willing to plan an attack for days or for years?). It is important that this adversary description be grounded in reality, and that it includes any adversary groups that are known to harbor hostilities to your specific facility or to other similar facilities. Many resources are available to help you understand the threat environment for your system, including local and regional crime statistics, interviews with law enforcement personnel at all levels (including the FBI), and security trade groups and publications. Your local InfraGard chapter may be able to provide some contacts to help you get started.

Blended Cyber-Physical Attacks

Using the definitions above, it is apparent that a blended attack involves the use of both physical and cyber attack tactics in the same scenario. So, while the ultimate target of a physical attack may be to manipulate or damage a particular component, a blended attack pathway may use cyber attack tactics to manipulate or disable cyber-controlled elements of the physical protection system (e.g., detectors, alarm annunciators, or locks) in order to enable the physical attack to be accomplished more easily. This type can be called a cyber-enabled physical attack. Similarly, the ultimate target of a cyber attack may be to manipulate a system's cyber controls, but a blended attack pathway may use a physical attack to access cyber control or entry points (e.g., network terminals or control rooms) from which cyber attacks are then launched. This type can be called a physical-enabled cyber attack.

A key characteristic of a blended attack scenario is that both the cyber and physical attack portions of the scenario must be successful in order for the adversary to "win". We start by considering a two-stage attack that involves either a cyber attack followed by a physical attack (cyber-enabled physical attack) or a physical attack followed by a cyber attack (physical-enabled cyber attack). In each case, the first stage of the attack is intended to bypass important security features of the second-stage attack target. If either stage of the attack fails, the adversary still has the option to perform a single-stage attack in either the physical or cyber realm. Thus, we can defeat a blended attack by defeating either the cyber-attack stage or the physical-attack stage, but the adversary can still win if a single-stage attack is successful. The question for the adversary is really very simple: "Is my success likelihood from a two-stage attack improved enough compared to my single-stage attack options to justify the added risks and complexity?" The adversaries choose whatever attack path they find most advantageous. Consider the simple example of a 12-foot wall that contains a computer-controlled gate. Is it easier to have a large co-conspirator boost you over the wall, or to have a techno-savvy co-conspirator remotely hack into the gate's control system to open a gate? Our job as security risk analysts is to identify all of these potential problem scenarios, understand their risks, and provide remedies as appropriate. (Note: One can consider attacks that use more than two stages in a similar manner. Handling multiple stages is not necessary, however, until the set of possible two-stage attacks has been analyzed and is well understood.)

Consider another example of a cyber-enabled physical attack. We assume that an adversary who controls a particular cyber node controls all hardware components associated with that node. Thus, the adversary can turn off all detectors, silence all alarm annunciations, and defeat all barriers (e.g., unlock doors) that can be accessed from that cyber node. One can assess this hobbled physical security system by removing these now-defeated components (i.e., acting as though they no longer exist) and conducting a traditional physical security analysis to understand the conditional likelihood of adversary success, given that they control a particular cyber node. One can also conduct a traditional cyber security analysis for this target node to understand the adversary's success likelihood for the initial cyber attack in this two-stage blended attack. The blended attack is successful only if both attack stages are successful. So if the likelihood of success for both the cyber and physical attack steps are expressed as probabilities, they can be multiplied in a logical "AND" operation to find the overall adversary success likelihood for the blended attack. We compare this likelihood for blended attacks with adversary success likelihoods for single-stage physical and cyber attacks. The attack where the adversary has the greatest success likelihood is their optimal attack path, and should be given higher priority for remediation decisions. This process is repeated for various combinations of compromised cyber nodes until the cyber effect on physical security is well characterized. Note that the adversary's optimal path may change as different sets of cyber assets are compromised. [5]

An important challenge for the security risk analyst when evaluating blended attack pathways is to estimate the adversary success likelihood for a cyber attack. Techniques for evaluating cyber security systems are less developed than the techniques used to evaluate physical security systems. Sandia has developed an approach that compares the strengths and capabilities of an assumed adversary with the characteristics of the cyber security system to estimate the range of success likelihoods for that adversary based on evidence theory. [6, 7] This category-based approach evaluates adversary groups according to six attributes:

- ***Goal Commitment Intensity*** – Level of determination in achieving their goals or objectives.
- ***Stealth*** – The ability to achieve the level of stealth necessary to achieve the adversary's goal.
- ***Physical Access*** – Ability to gain physical access to the cyber system design information and/or a cyber resource that is needed for the attack.
- ***Cyber Skills*** – Sophistication of computer knowledge such as the ability to develop new attacks rather than merely re-use existing, known attacks.
- ***Implementation Time*** – The total amount of time that an organization is willing to use in planning, developing, and deploying an attack (i.e., organizational patience).
- ***Organization Size*** – The size, structure, and social networking ability of the cyber-literate members of the adversary group.

Funding can enhance adversary capabilities for each of these attributes. The characteristics of the cyber security system are measured in categories called "security primitives" that include authentication, access control, encryption, integrity checking, data-aging protection, monitoring, and system management. [6, 7] For example, authentication works to assure that a person is who they say they are, and is often implemented using passwords. Network access control works to limit access to systems via communication links. User access control determines the system privileges that each user has.

The range of attack success likelihoods is estimated by identifying plausible network attack paths, assessing security along that path according to the above security primitives, and then comparing the adversary's strengths and capabilities with the assessed cyber security characteristics using an expert-based tabular rule set. The goal is to understand the likelihood that a specified adversary category can traverse an attack path and thereby complete a successful attack. [6, 7] (Note: This evidence-based approach is similar to past work on attack trees and attack graphs [4, 8]. However, it better captures the uncertainty in the user inputs to the cyber/physical security assessment, and hence in the conditional risk estimates for the overall cyber/physical protection system.)

Results

Sandia National Laboratories has developed and published the method described above [6, 7, 9, 10, and 11]. As part of the development process, several example facilities were examined. The methodology is robust and produces important observations. For example, adversaries may pursue a physical-only attack for facilities with well-protected cyber assets. Different combinations of cyber protection, physical security component cyber control, and adversary cyber attack capabilities can yield very different optimal attack scenarios for the same facility. A business partner's poor cyber security can degrade another facility's physical security posture. Evidence-based techniques [5] for evaluating cyber security systems provide the security analyst a way to combine these insights with the details of the optimal scenario paths to pinpoint security system weaknesses and propose cost-effective solutions.

Minimizing Blended Attack Potential

The best way to avoid blended attacks is to do two things:

- Implement existing physical and cyber security best practices for your industry segment, and
- Understand how interactions between the physical and cyber components of your system interact in ways that could lead to vulnerabilities or consequences.

These techniques provide a way to understand the cyber-physical security interactions. Best practice lists for physical and cyber security are numerous and industry-specific. However, the authors have observed that these best practices are often implemented piecemeal, without balance. Companies often focus strongly on physical security while (comparatively speaking) neglecting cyber security, or vice versa. Cyber security is particularly easy to neglect because the effects of "deficient" cyber security are often less obvious than a missing fence or gate. The following ten questions can help you decide how best to control physical and cyber access and improve security:

1. Are there adequate controls on *physical* access to your most critical system components (e.g., power or commodity distribution) *and* cyber nodes (e.g., control rooms, servers)?
2. Are your users choosing strong passwords and changing them on a regular schedule?
3. Are your business and control networks interconnected to the least extent possible?
4. Is remote access to your business networks and control networks only enabled when absolutely necessary? Consider remote connections from both employees and vendors. A weak cyber security posture at one of your vendors (e.g., a hardware vendor with remote maintenance access) can negate your state-of-the-art physical security system.
5. Is two-factor authentication being used for remote access to your most critical applications?
6. Are user's access privileges reviewed (and possibly modified) when their job function changes? Are former employees' access privileges revoked promptly? These include physical access privileges (IDs, badges, keys, etc.) and cyber access privileges (network accounts, etc.).
7. Do you regularly test and audit your access control and intrusion detection systems? This includes systems for physical security (e.g., alarms, procedures, and entry systems) and cyber security (e.g., firewalls and other network access control devices).
8. Are your virus checkers, operating system patches, and spyware removal software current?
9. Do you regularly audit for unauthorized physical security conditions (e.g., unlocked doors) and cyber connections (e.g., dialup modems that bypass your site's access policies)?
10. Do you have up-to-date disaster recovery or mitigation plans that have been recently tested and rehearsed? These plans should include elements such as cyber backup procedures, mutual aid agreements, priority service or supply contracts, law enforcement contacts, and emergency power or communication systems.

When the policies and practices described in these questions are in place, the techniques described in this paper can help ensure that your site has a security posture that balances spending and protection between physical and cyber security elements. These techniques can also help balance between prevention options (e.g., cyber/physical security systems) and mitigation/repair options. This balance helps ensure that your security system is effective and cost-efficient against physical, cyber, and blended attacks.

Acknowledgments

The Laboratory-Directed Research and Development Program at Sandia National Laboratories supported this work. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000. SAND2007-xxxxx, Unclassified Unlimited Release.

References

1. H. A. Bennett, "The EASI Approach to Physical Security Evaluation," SAND76-0500, Sandia National Laboratories, Albuquerque, NM, 1977.

2. M. L. Garcia, *The Design and Evaluation of Physical Protection Systems*, Butterworth-Heinemann, 2001.

3. Systems and Network Attack Center (SNAC), National Security Agency (NSA), "The 60 Minute Network Security Guide (First Steps Towards a Secure Network Environment)", Updated July 12, 2002, http://www.nsa.gov/snac/support/sixty_minutes.pdf.

4. Amenaza Technologies, SecurITree software, http://www.amenaza.com/

5. J. Darby, "Evaluation of Risk from Acts of Terrorism: The Adversary/Defender Model Using Belief and Fuzzy Sets", September 2006, SAND2006-5777, Sandia National Laboratories, Albuquerque, New Mexico.

6. J. Darby, J. Phelan, P. Sholander, B. Smith, A. Walter and G. Wyss, "Evidence-Based Techniques for Evaluating Cyber Protection Systems for Critical Infrastructures", IEEE MILCOM 2006, October 2006.

7. J. DePoy, J. Phelan, P. Sholander, B. J. Smith, G.B. Varnado, G.D. Wyss, J. Darby, and A. Walter, "Critical Infrastructure Systems of Systems Assessment Methodology", Sandia National Laboratories, Albuquerque, NM, October, 2006, SAND2006-6399

8. R. P. Lippmann and K. W. Ingols, "An Annotated Review of Past Papers on Attack Graphs", Technical Report ESC-TR-2005-054, MIT Lincoln Laboratory, Lexington, MA, 2005.

9. K.A. Gordon and G.D. Wyss, "Comparison of Two Methods to Quantify Cyber and Physical Security Effectiveness," SAND2005-7177, Sandia National Laboratories, Albuquerque, NM, November 2005.

10. J. Darby, J. Phelan, P. Sholander, G.B. Varnado and G. Wyss, "A Cyber-Physical Security Assessment Methodology (CPSAM)," Presented at the American Nuclear Society Winter Meetings, November 2006, American Nuclear Society, LaGrange Park, IL, 2006.

11. J. Depoy, J. Phelan, P. Sholander, B. Smith, G.B. Varnado and G. Wyss, "Risk Assessment for Physical and Cyber Attacks on Critical Infrastructures", IEEE MILCOM 2005, October 2005.