SCADA Threat Modeling Using Attack Scenarios

Ralph Langner
Langner Communications AG
rl@langner.com

Bryan L. Singer
Wurldtech Security Technologies
BSinger@wurldtech.com

Abstract: The threat of cyber security failures in industrial automation and SCADA is ever present and even increasing, but there is often little understanding as to what the threat actually may be. While engineers and security professionals can often measure the vulnerabilities in a system, and can apply quantitative logic to the impact analysis, one of the key components to understanding risk often goes ignored. The challenge is that the threat is a key component of risk.

Understanding the threat is tantamount to creating awareness, without which a business case and justification for security expenditures cannot be developed. Most organizations exist happily in the thought that, "this could never happen to us," but to those informed about security, it is not a matter of if, but when. Without sufficient threat data, however, developing threat models and attack trees are a matter of academic interest with often low perceived value to the business community.

This paper takes an alternate approach to building threat data. Leveraging past known cases of cyber security failures as well as recent research and development into the areas of industrial cyber security, this paper adopts an approach of creating scenarios that explain the why, how, and consequences of various attack vectors and compromise scenarios. The paper begins with expressing the traditional based approaches to understanding risk, the various strengths and weaknesses in ascertaining risk, and then tours various threat scenarios that can be used or extended by the industrial automation and SCADA community to help create an accurate picture of risk and utilize in business case development for security countermeasures.

Keywords: Threat Modeling, Cyber Terrorism, Risk Assessment

1 No Threat, No Risk: Towards a Scenario-Based Risk Model

If there is a terra incognita in the SCADA security landscape, it's threat country. We know quite a lot about vulnerabilities, about impact, and we even have a good selection of countermeasures available. But when it comes to determining threat, we know little,

and little research is done in this area. This is also a critical missing element as understanding the threat is the key component to creating appropriate awareness upon which all decisions about justification for security improvements hinge. This is surprising because the general importance of threat assessment is well known among security experts:

> "Threat assessment is probably the most important part of the entire risk assessment process." (Carl A. Roper, Risk Assessment for Security Professionals [1])
>
> "Who are the attackers? What do they want? What tools are at their disposal? Without a basic understanding of these things, you can't reasonably discuss how secure anything is." (Bruce Schneier, Secrets and Lies [2])
>
> "The security of an application cannot be understood, analyzed, or characterized unless the threats to the system are quantified in a threat model document." (Frank Swiderski & Window Snyder, Threat Modeling [3])

The little insight that we have on SCADA threats is that the research community seems to identify intentional malicious compromising by "hackers" or terrorists as the number one threat for SCADA installations, especially for those in critical infrastructure. Nearly half of all documents about SCADA security on the Internet mention attacks. On the other hand, there is almost zero empirical evidence to support this theory, which raises the question of threat credibility. Here's the problem. It is commonly assumed that

$$\text{Risk} = \text{Threat} \times \text{Vulnerability} \times \text{Damage}$$

The people who invented this equation did not mean it metaphorically but in a strict arithmetic sense, and it is used this way every day in the insurance industry. It can, and is, also applied to information security. Threat in this equation is the number of empirical threat manifestations for a given timeframe, such as the number of hurricanes per year in South Florida, or the number of computer viruses and worms per day on the Internet). No matter how we put it, without past damage history, the threat coefficient for targeted SCADA attacks is close to zero in this model. Since the multiplication of any value by zero returns zero, the resulting risk also is zero. So actually there should be no reason to be concerned about SCADA attacks, and as a matter of fact, this reflects the basic thinking of decision makers. It may be discomforting for some, but if we apply common wisdom, it looks like SCADA attacks are the emperor's new clothes of the SCADA security community.

A converse approach promoted by many in the US Government and in SANDIA's RAMCAP methodology is to set the threat equal to one and focus on vulnerability and damage. This is certainly the most conservative approach to assume an omnipotent threat to be always present waiting for a vulnerability to inflict damage, especially with critical infrastructure assets where the damage could be substantial in loss of life and massive damage to an economy. Experience has shown that this approach has not been persuasive to date to senior management who needs to approve money and other resources to address risks they are convinced exist.

Given empirical attack data, quantifiable risk calculations can provide great information to management in terms of potential economic impact, but if the threat never materialized, such numbers may be unrealistic and even misleading. Based on empirical

incident data, the risk of a terrorist attack by passenger airliners flown into skyscrapers calculated to zero before September 11, 2001, but obviously the risk did exist. The reason why we are dealing with the concept of risk is to prevent damage or to reduce the amount of anticipated damage. Our understanding of risk should be a reasonable guidance for countermeasures. If this cannot be done in the classic risk equation, it's time to use another model. For situations with no past incident history but other indications of risk, the authors suggest a scenario-based risk model. This model uses the same components of the traditional risk equation. The difference is that the link between the risk parameters is not arithmetic but logic. In order to do that, we put the elements of a successful attack into a logical order of cause and consequence and arrive at the following axioms:

1. An attacker , or threat, who successfully exploits vulnerabilities of a given target will inflict damage. Or, graphically:

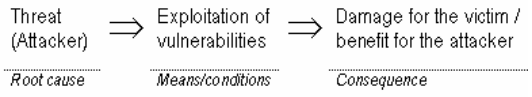

2. Risk is a real-world combination, a "scenario", of the elements from 1 where the conditions and requirements for success are met. With all conditions and requirements for a successful attack being met, it becomes a matter of stochastics if an attack is launched; depending on factors like potential attackers being aware of the target and its vulnerabilities.

3. The relevance of a specific risk for a given target is defined by the matching of the risk scenario with the target environment. If certain elements of the scenario – such as the use of confidential production processes that present a high market value – do not apply for the target environment, the specific risk is irrelevant for this target and may be ruled out. Some scenarios just don't fit, just like the risk of getting breast cancer for male persons, or the risk for a nuclear reactor to be blown up by a teenage hacker via the Internet. In practice, much more risk scenarios can be ruled out for a particular target than be identified as relevant – which is simply a reflection of the low incident rate.

4. The rationale for risk mitigation and the efforts/resources that should reasonably be put in it result from risk relevance and from the anticipated amount of damage.

While 1 is almost self-evident, the real starting point for threat modeling as proposed in this paper is 2: the conditions and requirements for a successful attack. The same basic concept can actually be found in Schneier's attack trees [4], with the difference that these are about the technicalities of an attack on a micro-level, while here we are looking at potentially successful attacks on a macro-level. An attack tree is not constructed by using statistics of empirical incidents but by using educated reasoning about which actions could – and *would* – compromise a system. It is not about attack statistics, it is about attack requirements.

This approach is quite straightforward when applied to situations where we creatively fill in different elements based on the characteristics of the target environment. In a threat-centric approach, we may start with the goals of identified potential attackers such as terrorists, evaluate suitable attack strategies by applying the knowledge of professional penetration testers, and then determine the likelihood of an attack by analyzing how much of the attack requirements are met. In a target-centered approach, we may start by analyzing the characteristics of the target system, including vulnerabilities, and from there work up to identifying attackers who could benefit from compromising the system. As a result, it gets comparatively easy to determine the realism of certain scenarios, and even to furnish dedicated counterstrategies.

If this is compared to the traditional arithmetic risk equation, we see that the scenario-based approach doesn't result in metrics. The authors don't see this as a problem because we get plausibility in exchange. Basically, every risk is hypothetical; it is a prediction about damage to come. The difference here is that risk is not expressed as a number, such as a percent value, but as a realistic story. In fact, *every* risk is a fictional story. It tells us about a fictitious damage. If the damage was factual, it was an incident report. Of course, if the story's main character – the attacker – is characterized as some anonymous evil computer genius who is able to break into major facilities just like magic and mostly for recreational purposes, the story is as realistic as a fairy tale. If a risk scenario isn't plausible, it won't and shouldn't prompt action. As a general rule, plausibility of risk, and therefore acceptance for the necessity of risk management, can best be established by composing realistic risk scenarios.

2 A Threat Modeling Framework for SCADA Attacks

Threat modeling is helpful in any situation where we don't have empirical evidence from past documented incidents and still want to determine risk because the target system has many vulnerabilities and the impact of exploits could be huge. Threat modeling basically consists of collecting all the threats that we know about or that we can think of. For those threats that we don't have empirical evidence for, experienced security experts can put together hypothetical what-if-scenarios that present a risk for the target system.

If we limit our scope in this paper to intentional malicious acts[1] against SCADA systems the threat modeling process starts by asking the following questions:

- What are realistic anticipated attack goals? Given the characteristics of the target system, what purpose could compromising the system serve? Who is pursuing such a purpose?
- What would be required to successfully carry out such an attack?
- Given the capabilities of identifiable potential attackers, the cost and benefit of an attack, how likely is such an attack?

It takes little consideration to determine that a SCADA attack wouldn't "just happen". In the Western world, SCADA attacks are criminal acts if we exclude military scenarios.

[1] There are other threats that are not discussed in this paper, such as force majeure and malware.

Any SCADA attack that justifies its name is a complex operation that requires planning, if just to lower the risk of getting caught, maximizing damage and increasing the likelihood of achieving the attacker's goals. A criminal action is usually evaluated in terms of three things: Motive, opportunity, and means (MOM). Proving a criminal action typically requires demonstration that each of these are fulfilled. Risk analysis often tells us plenty about vulnerabilities (opportunities), and means (the how of exploiting a vulnerability), but a critical piece is often missing in understanding the motivation behind an attack. Understanding the potential motivation behind these exploits rounds out the risk picture and brings the true threats to control systems to light, and helps us clearly demonstrate the risk to management.

2.1 Attack Goals

The attack goals for a SCADA system have little in common with the attack goals of a typical office IT system. There are no credit card databases to download. There is no customer database to destroy. There are no dumb users in front of the terminal that could be lured into a phishing attack. You can't use a SCADA system for money laundering, and there is no child pornography on a SCADA server. Most of the usual things that conventional cyber criminals do on a daily basis don't work with SCADA systems. Most systems are not even connected to the Internet. For the average cyber attacker, a SCADA system has no value.

The most notable characteristic of a SCADA system is that it controls or monitors a physical or chemical process. Thus the process will be the utmost target of the attack. As such, SCADA attacks are a class of its own. You would rob a bank because that's where the money is, and you would compromise a SCADA system because that's where process control is. Any attacker not interested in process control will find much better targets than a SCADA system.

Every attack on a SCADA system will either manipulate the target, including destroying the target, presumably in order to manipulate or stop the physical or chemical process controlled by the system, or steal something from the target, such as recipes or other intellectual property of significant value. The few things that can reasonably be achieved with a SCADA attack can be used to group attacks into the following clusters:

- Sabotage. Every targeted attack intending to manipulate (deny, disrupt, or deter) process control is summarized in this category. Attackers can be subdivided into insiders, hacktivists, and organized crime. The threat of sabotage can also be used for extortion.
- Cyber terrorism. Cyber terrorism is actually a special case of sabotage that we put in a category of its own because the ultimate goal of process manipulation is the achievable collateral damage, ranging from impaired social and economic life to injuries and death of innocent people. The result of a cyberterrorist attack is not limited to a specific target facility.
- Espionage. In an espionage scenario, the attacker intends to steal confidential information from a SCADA system – most likely recipes or other secret

product / production details that cannot be more easily be obtained by other means.

2.2 Attack Requirements: Capabilities

Everyone has the capability and access required for shop lifting. You just walk right into a shop, pick the item you are interested in strongly enough, and leave without paying. For SCADA attacks, this is different. Special capabilities are required, and access to virtually every installation is restricted. A typical SCADA installation may have numerous vulnerabilities, but they certainly cannot be exploited by everybody, even if a strong intent to blow up a facility is asserted. A SCADA attacker must be knowledgeable in multiple fields:

- Generic IT knowledge: Knowledge of TCP/IP-based networks and on classic attack methods and tools (WLAN, scanners, sniffers, denial of service (DoS) etc.). The general knowledge that attackers of conventional office networks and environments use can to certain extent also be leveraged against plant networks. Programming skills may be required for more sophisticated attacks.
- SCADA knowledge: IT knowledge on SCADA, industrial networks, DCS, PLCs, protocols, and on popular products. It turns out that typical SCADA installations show many vulnerabilities that are far more easy to exploit than Web and enterprise applications.
- Industry or process specific knowledge: The general technical knowledge about SCADA may not be fully leveraged until process knowledge comes in. This makes an attack much more powerful as the specific physics, chemistry or logistics of the process may be targeted to increase damage or "hit where it hurts most".
- Insider knowledge: Site specific intelligence. The attacker knows the what, when and where of the target facility. This knowledge can be used for very targeted attacks, and to multiply damage even more than just with process specialist knowledge. A malicious insider may even go about determining vulnerable assets with cascade effects that the asset owner failed to explore and secure systematically, such as seemingly unsophisticated support systems (vapor, energy, water, heating, cooling, pressure etc.) that affect multiple, complex procedures. -- On the other hand, employing site specific insider knowledge limits the suspects to employees, contractors or other insiders.

Certain capabilities can also be "outsourced". For example, insiders may be bribed, malware can be purchased, and experts may be contracted. Money and other resources, such as religious or political influence, may be used to acquire the capabilities needed for an attack. The easiest capability that can be achieved this way is expert knowledge. As experience shows, there is already a marketplace for exploits, and for custom designed malware. It would be naive to assume that this market is limited to office IT attack vectors. All this suggests that the commonly used concept of “security by obscurity” is bound to fail.

2.3 Attack Requirements: Access

A system cannot be attacked unless the adversary has access to it. Any attack must sooner or later hit its target physically or logically, which also involves some kind of physical connection. While it is often implied that cyber attacks may be launched via the Internet, this is highly unlikely as the majority of SCADA systems are not connected to the Internet and an online attack could easily be recognized. Therefore, for a targeted attack, the attacker usually must be able to access the target by other means, often by physical access to the premises. For several attack scenarios this imposes a big hurdle, as an attacker with enough motivation and capability still is unable to mount the attack if he has no access to the target. Gaining access to the target can be the most difficult task for the attacker, and especially the most dangerous because of the risk of getting caught.

Access is only loosely coupled with capability. For example, unskilled workers may have access to SCADA systems, but they lack the capabilities required to mount a cyber attack against those systems. Capability alone is not sufficient to launch an attack, and neither is access.

While it can be argued that access requires capability, the capability for gathering intelligence for example, there are still reasons to put access in a requirements category of its own. Access can be controlled, but the acquisition of capabilities usually cannot.

2.4 Attack Moderating Factors: Opportunity and Constraints

Other factors that influence the likelihood of an attack, the target selection and the timing are opportunity and constraints:

- Opportunity. Opportunity is capability and/or access that, just by chance, comes for free. Opportunity is nothing that an attacker "has". There are attacks that are completely opportunistic, like the Pennsylvania wastewater hack, and other attacks that don't rely on opportunity at all. The latter can usually be found in military scenarios where several contingencies are provided in order to guarantee success for a critical mission.
- Constraints. Constraints are scenario-specific factors that limit the attacker's options. For example, an espionage attack might make sense only within a specific timeframe. A terrorist attack might make sense only if it can be associated with a "symbolic" target. Simply blowing up a small family-owned factory in the middle of nowhere could be regarded as a failure. For an espionage scenario, it is essential that the attack remains undetected. For sabotage scenarios not carried out by insiders, it is usually desired that somebody or even the whole world takes notice. A sabotage act that goes unnoticed could be a failure, same as an espionage attack that gets noticed.

Once the basic concept of an attack plan is set, constraints limit the choices of suitable attack methods, and opportunities may make some targets look more attractive than others. For example, if an attacker's goal is to fly passenger airplanes into skyscrapers using nothing but box cutters and fake bombs, it follows almost naturally to pick transcontinental flights where the fuel load is high and the number of passengers is low.

3 Attack Scenarios by Category

In this section we will apply the threat modeling framework to compose and discuss several sample attack scenarios for different attack categories. It is not implied that every conceivable or even every likely scenario is captured.

3.1 Sabotage by an Insider

Though not happening on a daily basis, conventional sabotage acts by insiders, such as attacks not involving control systems and automation equipment, are well known. Conventional sabotage acts are probably as old as industrialization and will presumably have to be considered a risk as long as the environment is not operated 100% automatically. The motivation for sabotage by insider is either revenge for perceived setbacks or the simple desire for extra breaks. Suspects are disgruntled employees and contractors.

The problem with sabotage by insider is the high level of insider knowledge and the access to systems. Insiders usually have both capability and access sufficient to create major disaster for the target. On the other hand, the number of suspects is limited. The biggest constraint for the insider attacker is the risk of getting caught, as few people are willing to risk losing their job and probably even go to jail just for expressing frustration or enjoying an extra hour of lunchtime. The risk of getting caught gets even bigger if special capabilities such as programming skills are utilized. For example, the ability to manipulate database tables plus the access to those systems "where it hurts most" is usually very limited. From the few people in the organization that match the requirements, only one or two may have exhibited the frustration and criminal energy to qualify as a suspect. Therefore, sophisticated cyber attacks are much less likely than raw hardware attacks on SCADA systems and automation peripherals. A reasonable obscuration tactic is to use attack methods that "could have been done by anyone".

3.1.1 Sample Attack Scenario: Basic Physical Attack

Much of automated process control, and therefore of the smooth operation of a complex facility, can be disturbed by simple physical manipulations "at the right place". Examples include rewiring an Ethernet patchboard, pouring a glass of water onto a PLC/RTU with critical functionality and resulting in cascading effects, or cutting power of a SCADA/DCS server.

In all cases, halted production for up to several hours can be expected. The most critical components are those that are accessible by many, as the risk of getting caught decreases. If the target device is accessible only by one or two, the risk of getting caught is much too high.

Advisable countermeasures are comparatively simple. If the physical access to IT systems and automation peripherals is restricted, a suspect either cannot mount the attack or will experience the risk of getting caught as prohibitive. Restricting physical access by putting systems into locked cabinets will usually be sufficient, and is not associated with high cost.

3.1.2 Sample Attack Scenario: Software-Based Attack

Networked automation peripherals create vulnerabilities by themselves, as they enable process manipulations without touching a SCADA system. Every PC in the network may be used to launch a software-based attack which is almost impossible to be identified as such and backtracked unless specialized industrial versions of intrusion detection systems or forensics tools are in place.

It should be noted that little to no programming skills are required to launch this attack. Several HMI, OPC and other software demo versions can be downloaded from the Internet for free, and in many cases even full versions exist in their cracked forms on warez sites, the only need to be installed and configured properly in order to be abused for unauthorized process manipulation in a real production environment. What the attacker does need, however, is specialized knowledge of control system protocols and products. Detailed process knowledge may also be required to make sure that the attack does not result in injuries of co-workers, which are generally not desired by saboteurs. Such information can be obtained either through Internet researchers, or could even be coerced or "social engineered" out of plant personnel. There have been a number of known cases where social engineering has been utilized to gain process information, such as environmentalists, or disgruntled employees have wilfully given such information away. It can also be expected that the attacker has sufficient knowledge about the IT infrastructure to evade system and network monitoring or to be able to sufficiently cover ones tracks once a compromise has been executed.

The basic countermeasure against this type of attack is to limit access to PC's that may act as attack launch pads. The attacker will likely not use "his" dedicated workstation, in an attempt to misguide forensics. Since this attack requires installation of third party software, it is helpful to prevent this possibility for easily accessible systems by using monitor/keyboard terminal solutions and placing the "real" system in a locked room. Additionally, it is often a good measure to block access from the office networks to any direct industrial protocols and leverage other technologies such as terminal services to prevent unauthorized machines from the Internet or office networks from communicating directly with process control assets. Other than that, wireless access should be limited and/or tightly protected. In certain industries such as wastewater, attacks of the Vitek Boden style using radio modem to compromise SCADA systems and automation peripherals in remote facilities would still succeed today as there is little to no encryption and authentication provided by standard products.

3.2 Sabotage By Outsider: Organized Crime

An attacker may extort money from the target organization by threatening to manipulate or halt the production process. The capability is usually demonstrated when demands are made to make the threat credible.

As will be indicated below, an extortion scenario is costly for the attacker. It will pay off only if he is able to cash in big. For the victim, on the other hand, the demand must be considerably lower than the anticipated damage, which must be devastating for the target. The attacker will likely choose targets where the initial damage, such as manipulating or halting a production process, producing deliberate environmental

damage, will trigger a follow-up / cascade damage for the target, such as bad press or governmental restrictions. Therefore, it is likely that the target will be chosen in specific industries such as pharmaceutical, chemical, food & beverage (baby food is more likely than pet food), consumer products such as cosmetics. Power utilities may also be considered. Other qualifying targets can be found in industries where uninterrupted production is essential and production disruptions of less than an hour can be devastating. Examples can be found in the automotive industry with its tightly coupled just-in-time/just-in-sequence logistics chains.

The attacker must meet a comparatively big set of requirements. Actually, he must have all four types of knowledge mentioned in the previous section. He must have generic IT knowledge in order to mount an attack that is detected only after it's too late. He must have control system knowledge in order to manipulate process control. He must have industry or process specific knowledge to determine where he can hit the victim significantly. He must have insider knowledge to determine exactly where and when to hit best. Besides that, he must have set up a plan for collecting payment with minimum chance of getting caught in the process.

The biggest constraint for this type of attack is the need to demonstrate attack capability when demands are made, and to convince the victim of second strike capability if demands are not fulfilled. After the threat is announced, the victim will raise security to the maximum, presumably with help from professional crime fighters and law enforcement.

3.2.1 Sample Attack Scenario

The attacker or group of attackers has decent insider knowledge from several years of workplace experience. The specifics of the chemistry, biology or physics of the target's industry are well understood, and also how the process can be manipulated with little chance for detection. Besides all that, the attacker is well aware of the victim's contractors with almost unlimited access to critical systems.

The attack is carried out via a contractor. Contractors which are basically IT companies with small budget and virtually no security are an easy means to get to the heart of critical process control. Once that the attacker has managed to infiltrate the contractor's network, he can use remote access lines to the target facility or compromise the contractor's notebooks that will sooner or later be connected to the target production network for maintenance purposes. Potential victims include all the organizations that the contractor works for, which are often within the same industry and using similar or identical SCADA/DCS products.

The actual attack is not carried out online, as it can be expected that the victim will disconnect all outside links that he is aware of once that the attack is detected. The attacker is using a custom designed malware that executes a software-based attack against control systems by manipulating specific process variables, or by simply bringing the process to a halt with no chance for quick recovery. Since the malware is custom designed, it is not detected by antivirus solutions that the victim may have in place. The malware is installed on multiple systems, ideally in multiple facilities of the victim that are networked among each other. Self-replicating techniques may be used for the

distribution. The threat now is made that the second strike will be executed on other systems/facilities within a timeframe, such as 24 hours, that is long enough for delivering payment but too short to prevent the second strike.

Counterstrategies against this scenario start with controlling the access from contractors. Especially in situations where a good relationship with contractors is established over a long period of time, let's say for more than a decade, controls are viewed to be unnecessary and probably even as an indication of distrust. On the other hand it tends to be forgotten that a contractor technically extends the target's perimeter and therefore must be controlled the same way that access to local systems is controlled. Countermeasures include strong policies for local and remote access, controlling access with firewalls, and monitoring access. A second countermeasure is to restrict network traffic within the organization by network segmentation (i.e. firewalls). In a properly segmented network, the chance for spreading malicious code to other systems or facilities is much slimmer than in the "open" network architectures that even multinational corporations use. A third and very powerful countermeasure is host-based intrusion detection. Host-based intrusion detection solutions (HIDS) can be used to detect custom designed malware and thereby prevent attacks. Unfortunately, few organizations have installed HIDS on their plant floor systems.

3.3 Sabotage by Outsider: Hacktivists

Political activists are a well known threat to any organization that happens to get into their way and looks like a rewarding target to make their agenda public. Be it environmental pollution, employment "exploitation" of low wage third world workforce, usage of gene food, or fabrication of furs, there are all kinds of reasons why an organization can become a target for activists.

Activists are motivated by a cause, such as environmental or anti-globalization. One problem with activists is that they don't fear litigation, as they see it as some kind of martyrdom for a higher cause, and even as part of their PR campaign. Another problem is that many people that are attracted by such ideas are highly educated and skilled, and therefore in a good position to carry out a cyber attack. Once this happens, we are speaking of hacktivists who employ cyber attack techniques, a.k.a. hacking, in a context of political activism. Sabotage by hacktivists is basically extortion as they will usually intend to force the target organization into specific action, such as giving up certain procedures or even to ultimately shut down a specific plant. It is not hard to guess that if they had the capability for a cyber attack, they would use it. It may be expected that people bold enough to overcome conventional physical security measures and occupy the top of a smoke stack, for example, will sooner or later get out their notebooks to cause trouble in a more efficient way. Even if hacktivists had the capability to do so, they would presumably not blow up the target facility as this could result in loss of any credibility and get the attackers into jail for a long time. The major constraint for this type of attack is its publicity. The attack is intended to be public, and to be perceived as kind of fair, courageous and intelligent. In a David vs. Goliath setting, the weakness of the target organization has to be demonstrated, thereby destroying public confidence in the target.

3.3.1 Sample Attack Scenario

As the term suggests, hacktivists are activists that usually do not follow a long-term strategy in their operations. Their objectives will be either to disrupt operations or to gain useful information and then publicize the event to create a "smear campaign" against the organization or create other negative public situations. The objective is to make the target organization be "at fault" and to be able to claim credit for the exposure. As such, the attackers will not want to be associated as having created the negative problems, but rather as "exposing" them. For this reason, likely attack methods will focus on online hacking, given the opportunity. The relative low risk of detection and the ability to create damage without attribution to the attack provides a perfect platform to reach both of the above objectives. This could either be an unprotected WLAN access point, an open modem line or some ill protected Internet access.

The most likely attack strategy is either a DoS attack, using readily available attack tools, theft of potentially useful information in extortion or public smear campaigns, or a demonstrated manipulation of the production process that is then made public. In these cases, often the threat can be greater than execution, and simply obtaining enough credible information can be sufficient to create headaches for the target organization. Industry or process specific know-how for the latter may be available as some non-governmental organization (NGO) protesters go to great length to study details of their opponents. Some attackers may even be disgruntled insiders. Low-tech attacks are often combined with high-tech attacks as well. In known previous cases, social engineering is utilized to obtain sufficient information about a process to conduct some type of electronic attacks as well.

Any organization will usually know well in advance if it is the target for NGO protesters, but this may not always be the case. Organizations would be well served to consider their usage of hazardous chemicals, potential chemical or industrial waste issues, and even such issues as "eye-sore to the community." If this is an issue, countermeasures will make sure that no access to the organization's IT systems and network is possible, be it by WLAN, Bluetooth, or Internet. It should be taken into consideration that organizations cooperating/contracting with highly visible targets can also become a target.

3.4 Cyber terrorism

Cyber terrorism is a special case of sabotage by outsiders. The discriminating characteristic is that the terrorist ultimately intends to achieve a high level of collateral damage, accepting or even aiming at injuries and death of seemingly uninvolved civilians.

While it may be obvious to start looking for a terrorist threat in the Arabic world, domestic terrorism of the Theodore Kaczynski, a.k.a. Unabomber, style should not be neglected. As a matter of fact, SCADA cyber attacks against critical infrastructure organizations and against highly automated production processes in the private sector could be directly justified by referring to Kaczynski's manifest "Industrial society and its future". Kaczynski's stated goal was to get rid of industrialization and technology in favour of a return to "wild nature". Sample quotes from his manifest: "(...) it is certain

that technology is creating for human beings a new physical and social environment radically different from the spectrum of environments to which natural selection has adapted the human race physically and psychologically. If man is not adjusted to this new environment by being artificially re-engineered, then he will be adapted to it through a long and painful process of natural selection. (...) It would be better to dump the whole stinking system and take the consequences." – "(...) the two main tasks for the present are to promote social stress and instability in industrial society and to develop and propagate an ideology that opposes technology and the industrial system. When the system becomes sufficiently stressed and unstable, a revolution against technology may be possible." – "Until the industrial system has been thoroughly wrecked, the destruction of that system must be the revolutionaries' *only* goal." – " It would be hopeless for revolutionaries to try to attack the system without using *some* modern technology." [5]

Islamist terrorism should certainly be taken into account, too. While cyber attacks may look too sophisticated at first sight for organizations that favor bomb attacks as the premier attack method, it is a known fact that Arabic terrorists do have surprisingly high computer and programming skills, as documented by their use of the Internet for propaganda and communication purposes. As Richard Clark, former chairman of the President's critical infrastructure protection board pointed out, Al-Qaida was using the Internet to do at least reconnaissance of American utilities and American facilities [6]. Indications are that Al-Qaida is actively installing botnets for propaganda purposes. While many previous threats of "cyber jihad" have been met with little to no action, and it is unlikely they can mount a widespread attack on the Internet, using the Internet for propaganda distribution can certainly have tremendous effect. Further, these communications channels can be used to coordinate attacks, recruit more personnel, and distribute their brand of ideology. If we see Al-Qaida as some kind of holding organization for the management and financing of terror, it may well be expected that this organization is concerned about return on investment. Sooner or later somebody in the organization may determine that it could be a good idea to invest a hundred thousand dollars or so in a cyber attack against US critical infrastructure facilities rather than, or in addition to, sacrificing one of their followers to blow up two or three otherwise unsuspicious US citizens. As a matter of fact, published Al-Qaida statements indicate evidence for such plans [6]. It should not be a big problem for this organization to acquire the capabilities required for a SCADA attack. Although it is unclear if terrorist organizations do already have the capability to carry out a cyber attack, they certainly command the minimal funding required to buy such skills and products.

Another threat arises from Islamist states, most notably Iran, that are known for supporting terrorist groups, such as Hezbollah and Hamas, to pursue their goals. A cyber terrorist attack, organized and administered by a state-sponsored intelligence organization and carried out by a franchise terrorist cell, would be a premium example of asymmetric warfare and thus typical for states that cannot yet attack openly using military force. One might speculate though that a state-organized attack would target Israel's critical infrastructure, and it looks like the Israeli government is already aware of this threat and preparing for it.

A problem with all suspect groups mentioned is that they don't operate under a tight timeline. Whether the attack is carried out this month or in three years does not make a difference for the attacker, so there is ample time for preparation. Special problems with Islamist terrorists include the growing network of supporters and the non-concern about getting caught or even dying in the attack.

3.4.1 Sample Attack Scenario: Malware Attack by Domestic Terrorist

As the reference to the Unabomber has made clear, the possibility of a highly intelligent and educated terrorist with an anti-civilization and anti-technology agenda must be considered. Kaczynski's career culminated as a professor for mathematics at the University of Berkeley. His IQ allegedly ranged around 170, i.e. among the upper 1% of the most intelligent people. If we remember that Jeffrey Lee Parson, an 18 year old and not remarkably bright school boy was able to create the disastrous W.32 Blaster worm with little effort, and was dumb enough to incorporate the URL of his personal Web site in the code. It is easy to see that an attacker of Kaczynski's format would have no problem at all setting up a sophisticated, extremely aggressive cyber attack against SCADA systems.

The most likely strategy for this scenario is to design and implement custom malware that manipulates process control by attacking the top three or five SCADA, DCS and PLC product lines on the market. Widely used technology such as Modbus or OPC may be included in the attacks. Tracking statistics on security sites, we do see that popular protocol ports for Modbus TCP, DNP3, and others are frequently scanned and targeted, though no coordinated attack has been known to have occurred to date [7]. As the goal is simply to bring down the target facilities by creating a maximum amount of damage, accepting injuries and death, such a piece of software can actually easier be implemented than the scenario described in 3.2.1. An intelligent person that decides to dedicate most of his spare time to this task may end up with a very destructive "software weapon" within a couple of months.

Distribution of the "weapon" is achieved by conventional means via the Internet. For example, the attacker may use a Trojan horse approach, disguising the malware as some kind of freeware HMI or productivity tool that attracts the attention of operators who download it from the Internet, install and run it on their machines, and even pass it along to friends and business contacts in other organizations. Other conventional attack vectors and tools, such as Webattacker, or fake emails from major vendors informing about the need to install an urgent patch for security reasons, may be used. If done properly, victims may never figure out the cause of their problems, or only after significant damage has been done.

Unfortunately, there are no major constraints that would apply to this type of attack. The whole concept of the terrorist's goal and how to reach it may be as bizarre as transforming society to a state of "wild nature" by killing unsuspecting people with letter bombs, but as the Unabomber case illustrates, it may still materialize. An additional problem is that the resources that the terrorist is willing to put into the project don't have to match any reasonable return, as in scenario 3.2.1.

Advisable countermeasures include the general protection mechanisms against malware, with the understanding that conventional anti-virus solutions will be unable to detect the attacking software until a signature is published. Host-based intrusion detection will be able to detect the rogue software until the operator decides that the installation is ok and falls into the Trojan horse trap.

3.4.2 Sample Attack Scenario: Targeted, Supportive Attack by Islamist Terrorist Group

This scenario has already been discussed by various researchers. It assumes that a terrorist group attacks a critical infrastructure facility basically by conventional methods such as bombs. An attack of the facility's SCADA systems is carried out to support the conventional attack. For example, the SCADA attack may disable alarms or provide for cascading effects of the main conventional attack. Fire alarm sensors that are connecting to a master alarm server by Modbus/TCP, as seen on some facilities, may easily be compromised by launching a man-in-the-middle attack against the server, which continues to receive unsuspicious data while the facility is going up in smoke.

The SCADA attack is launched by an insider in the target facility. Such an insider may be recruited using established terrorist support networks. For example, Nidal Ayyad, one of the 1993 World Trade Center bombers, became a naturalized US citizen, graduated from Rutgers University, and worked as a chemical engineer at Allied Signal. Instead of using company stationery to order chemical ingredients to make the bomb, which Ayyad did, he could also have mounted a SCADA attack against Allied Signal's facilities. Such an attack can be physical, for example physically destroying SCADA servers as in 3.1.1 scenario, or be software-based. In this scenario, which assumes a bigger plot of an organized terrorist group, the actual SCADA attacker on site does not necessarily require any software skills. Attack software may be provided by some other member of the group that does not appear on the scene. The critical part for the insider is mainly to provide intelligence on the what, where, and when, and to provide access to the target systems.

Advisable countermeasures against this type of attack are basically identical to those mentioned in scenario 3.1.1. The focus is on preventing access to critical systems for unauthorized personnel. In addition to that, background checks of staff members in key positions can probably help to identify problems. It can be speculated that such an attack is more likely to happen in Europe than in the US, as US critical infrastructure organizations are generally better aware of such a scenario, whereas in Europe people who are potential supporters of terrorists are even less visible, and background checks are regarded as discriminating by some.

3.4.3 Sample Attack Scenario: Widespread Simultaneous Malware Attack

This scenario aims at crippling the economy of highly industrialized Western nations. The goal is not casualties in the first place but a substantial economic damage, combined with the message that Islamist terrorists are up to par on the use of high technology. Destroying Western economy is one of the outspoken goals of Islamist leaders such as

Omar Abdul Rahman, a.k.a. "the blind sheikh", who is serving a life long prison sentence in Florence, Colorado.

The assault weapon for this scenario is custom designed malware similar to those in scenario 3.4.1. However the emphasis is on brute-force DoS rather than on sophisticated process manipulations. The first attack stage is to create SCADA-specific malware that performs DoS on the most prevalent SCADA systems (the top 3 are sufficient), on OPC installations, and raw DoS and/or manipulation attacks against Modbus devices. This task is outsourced to people with the required generic IT and SCADA specific knowledge. In the second stage, which acts as a seed stage, the malware is distributed to key targets such as critical infrastructure organizations and big manufacturing companies with highly automated environments. Distribution is done by insiders. All the insider has to do is insert an USB stick or CD into one or more PCs on the plant floor. The insider does not need any knowledge about malware or about SCADA, he just needs access. The insider may receive the assault software at home from a Web site or by email, utilizing existing fundamentalist support networks. The malware is programmed as a logical time bomb. It gets into the assault stage at some point in the future when sufficient distribution can be expected and the chances for enhanced damage are high. A likely trigger would be Christmas eve, 2008. In the assault stage, the malware replicates itself using conventional worm technology, thereby infecting other systems in the facility and in other networked facility of the target organization. Thereafter, the malicious routines are executed. SCADA/DCS applications are knocked out, OPC servers are DoSed, automation peripherals are flooded with random writes. As a result, major facilities in the US and Europe will simultaneously be shut down, with cascading effects for their logistic chains. The effect to the general public may be significant, even if no casualties happen, and all caused by some piece of software that can probably be purchased for a five digit dollar figure.

For this scenario, we should assume that the malware is already in place. In the absence of a published signature, it was not detected by antivirus solutions. Therefore, counterstrategies involve system hardening and a well firewalled network to prevent the replication of the malware and unauthorized access of systems throughout the network. Certainly one thing that an asset owner should have is an emergency plan for such a scenario. One discomforting thought in this context is that Islamist terrorists are not in a hurry. Time works for them; with every year, they acquire more skills and more insider knowledge, and get more loyal followers in strategic positions.

3.5 Espionage

It can safely be assumed that acts of industrial and economic espionage take place on a daily basis around the world. If a company can invest just a few thousand or tens of thousands in espionage to take advantage of billions in research by another organization, then the motivation to conduct such acts becomes clear. The question here is in which respect process control can be considered as a valuable target. This is most likely the case where intellectual property is a critical asset of the production process, such as recipes or breakthrough production procedures that are not publicly known. However, a common misunderstanding is that espionage would be a concern only for high tech companies. There are several low tech industries where valuable procedures or product

designs may easily be obtained by screenshots of a SCADA application. Recorded attacks have been seen even at the small local competitor level, where an organization used information gleaned about how well their competition was performing for a given day to affect their own local pricing and thus compete more effectively.

Suspects for espionage attacks include local and foreign competitors as well as state-sponsored intelligence organizations. Attackers will, for example, be found in China in its struggle to become a, if not the, leading economy of the world. Besides, several intelligence services, especially in Russia, Ukraine and other Eastern Europe countries diversified to industrial espionage after the cold war and actively collect intelligence which is then passed to native end-users.

For an espionage attack, the attacker must be an expert. He must know what he is looking for, where to find it, and how to deal with countermeasures. He must master capabilities for getting access to the target, most likely including social engineering capabilities. A common constraint for espionage attacks is that the attack must go unnoticed, at least until the information in question has been obtained. Once that the attack is detected, the information can no longer be used as it may be purposefully compromised by the victim.

3.5.1 Sample Attack Scenario: Cyber Attack by Local Competitor

Local competitors in tight markets certainly have reason to conduct espionage. Whether the motivation is to win a new construction job, sell more of their own product, or just make a little extra for a short time, it is clear that there is monetary gain associated with espionage at this level. In this case, proprietary information such as recipes rarely is a factor. Consider a case where local competitors are both supplying cement to a series of construction projects in a geographic area. The main factors to influence the market for buyers is price, capacity to deliver, and quality. One of the two companies learned from their supplier that their competitor had a similar historian and laboratory systems. Using this knowledge, they are able to gain access through modem lines and using the same type of remote connection that the supplier used to access plant floor historians, quality systems, and others. Using only passive means of eavesdropping, they use the information gained in helping to increase their competitive position:

- When they see that the compromised company is having a particularly bad production day, for example outages, reclaimed product, etc. in the plant historian, they raise their prices to take advantage of what may be a temporary market condition.
- When they see that the compromised company is doing particularly well, such as first pass acceptable yield of product is high, they lower prices to influence buyers.
- If they detect potential quality problems by looking at the target company Laboratory Information Management System (LIMS), they on the sly pass this information to their customers that they "hear" there is a quality problem and thus should buy from the attacking company.

- If they were to gain access to ERP systems, the attacking company could use this information to understand buyer conditions as well, what orders have been won, pricing, etc.
- If the attacking company takes it a step further and actually compromises the process, one potential would be to modify production values, but also modify the lab samples. In cement, there are typically sample tests all the way out to 3-4 weeks from date of production. By this time, the product may well already be on the ground especially in ready-mix operations where the product is poured wet into a truck. If the problem is not detected until the 3-4 week lab tests, the company may end up in a costly recall situation.

In a commodity market or one in which price fluctuations are frequent, and in such a product where there is often very little variability in the actual product, if competition is stiff there are plenty of motivators to conduct such attacks. It is very common for integrators and other suppliers to use the same method all the time, often with the same passwords and users supporting the system, to "ease" the support of the operation. Contracts with all vendors should require confidentiality of operations when competitors are involved, implementing what is known as a Brewer-Nash or Chinese Wall methodology for internal procedures.

Further, despite the additional access control requirements, even vendors should be forced to change passwords on a periodic basis. Plant floor applications, such as historians and LIMS, are also often neglected in security hardening. Most view this information as relatively harmless because it takes someone with knowledge to understand the data. Hopefully this scenario shows that the need is not excused because of complexity of data; the often failed "security through obscurity" concept.

3.5.2 Sample Scenario: Combined Physical/Cyber Attack by Foreign Competitor

Espionage is carried out physically most of the time, involving physical entrance to the target facility. Once that the attacker has managed to do this, there are lots of ways to collect information that are much simpler than installing spyware on a computer system. A foreign competitor is mentioned in this scenario as it relies less on opportunity than scenario 3.5.1; it could, however, also involve a local competitor.

The attacker, who already managed to access the target's premises, simply physically steals a computer that stores valuable information. In many cases this is easier than it might sound. Disguised as maintenance staff from some contractor, the attacker disconnects the target system and walks out of the door with it. Although the theft may quickly be noticed, the victim may not be aware that it happened primarily for the information on the hard disk, not for the computer itself.

While on the premises, the attacker may also install a WLAN access point or radio modem that can then be used for cyber attacks. It is puzzling that such a "new" IT device rarely prompts employees for action. If employees recognize it at all, they assume that it must serve some legitimate purpose and fear that something might stop to work if

the device is unplugged. Approaching security with this issue may prove a false alarm or even look stupid.

Advisable countermeasures for this scenario include strong guard checks of equipment that leaves the premises, and regular checks of radio emissions.

3.5.3 Sample Attack Scenario: Cyber Attack by Foreign Intelligence Agency

Foreign intelligence agencies that acquire information for their home country's clientele are usually not targeting a specific organization, but a specific industry. The attacker therefore has a bigger freedom of choice in selecting a good target, and he may scan for favorable opportunities. Anyhow, the most likely targets of foreign intelligence agencies pursuing industrial espionage are not asset owners but small contractors who provide a high degree of knowledge. As outlined in scenario 3.2.1, it happens that such companies are also little guarded against cyber attacks and can thus be more easily attacked than a multi-million dollar organization that has some concept of protection against industrial espionage in place. After successfully breaking into a knowledgeable contractor, the attacker has it all: client data detailing confidential product and production specifics of multiple target organizations, and probably even online remote access to the respective live environments.

Such an attack may be carried out online via the Internet, as the targets in question almost certainly have Internet access and no air-gapped "production network" where the critical systems would be kept. With the tools and skills that the state-sponsored intelligence agencies command, network intrusions via the Internet have a good chance for success.

Advisable countermeasures are similar to those mentioned in scenario 3.2.1.

4 Unlikely Risk Scenarios and Lessons Learned

Every risk can be presented as a scenario. For SCADA attacks, this means to include information about the who, why and how of a potential attack, so that an evaluation is possible whether the major requirements for a given attack are met for a given target. A scenario-based approach to risk puts vulnerabilities, exploitation capabilities, access, and consequences (damage for the victim, benefit for the attacker, risk of getting caught) into context. This holistic approach makes it easier to assess the relevance of a given vulnerability or threat for a given target than an atomic approach that leaves it open who would actually use an exploit with any comprehensive reason to do so.

A vulnerability, or even an exploit, has zero relevance in itself for a given target. In the classic risk equation, even if vulnerability is 1 (or 100%), risk still calculates to zero if threat or damage is non-existent or unknown. Likewise, in a scenario-based approach, vulnerability alone is irrelevant, as the ultimate cause for damage, the threat, is non-existent or unknown, and if the consequence of a successful attack doesn't matter much for the victim. Unfortunately, most discussions of SCADA vulnerabilities don't extend to reasonable scenarios, or imply vague and unrealistic assumptions about attackers, attack goals, and potential damage. We can, however, evaluate the relevance of certain

commonly discussed vulnerabilities, exploits, suspect groups, and targets by putting together matching scenario elements like some kind of jigsaw puzzle. If scenario elements don't match, no picture emerges, and probably no risk is imminent.

4.1 Hacker Attacks

The term hacker commonly refers to a skilled person who attacks computer systems via the Internet mostly for recreational purposes. Without doubt, thousands of hackers go about their activity on the Internet every day. If we assume that hackers are motivated by technical challenge, SCADA installations with their outdated equipment would be among the least interesting targets for a hacker. Besides, most facilities are not accessible via the Internet. Internet access to a production network can only provide for opportunistic attacks. If a hacker has a choice to attack a top notch server farm or a production facility with five to ten rusty, low performing PCs, he will probably leave out the latter. Production networks are not among the most appealing targets for hackers.

Hackers are also not notorious for spending a lot of effort into vertical applications, so the average hacker wouldn't really know what to do with a compromised SCADA system at the application level. Once that he had gotten an idea of what he is on to, he would probably step back because of the anticipated, and presumably undesired, potential collateral damage. While people associated to the hacker community claim to be able to "hack" critical infrastructure facilities, it appears that such claims have largely been made as a demonstration of skills in order to fuel potential business for security consulting, which the respective people coincidentally happen to offer.

As has been pointed out, expert knowledge is a prominent requirement for targeted SCADA attacks. While most people, and especially executives, overestimate the effort required to become an expert in SCADA technology and products, the bigger problem for pulling off the more dangerous attacks is the industry and procedure specific knowledge that is required for leveraging the full power of the physics and chemistry of the target process. There are few people on the good side who really master IT security, SCADA, and process control at the engineering level. There are even fewer people with such knowledge on the evil side. As it turns out, the proverbial hacker is probably the type of attacker that any asset owner needs to worry about least.

4.2 Exploiting software bugs

Software bugs such as buffer overflows in SCADA applications and drivers, that may exploited to crash an application are hardly of interest for an attacker as long as a more dramatic effects may reliably be achieved by deliberate manipulations, including DoS attacks, of the network or automation peripherals. It is unlikely that such software bugs would play a role in targeted attacks. Implementation flaws are product and version specific. They are not something that an attacker can rely on because the bug is probably fixed a couple of days after the reconnaissance.

The custom-built vintage software typically found on many plant floor systems is full of bugs, but few people know about these bugs. Even if such software may reasonably be viewed as insecure, exploiting such vulnerabilities is even more insecure for an attacker. In a scenario-based approach, software bugs are more likely to pose a problem in

conjunction with random, unintentional threats, such as software or equipment updates, rather than in conjunction with cyber attacks.

Sometimes, vulnerabilities in SCADA systems are published that, if exploited, allow for the execution of arbitrary code on the target system. However, in some of the vulnerabilities discussed over the last couple of years, the exploit itself would require running malicious code inside the network in the first place. An attacker with the ability to use the exploit would therefore actually no longer need it. On the other hand, it has to be answered what the anticipated malicious code would do to a SCADA installation. As it has been pointed out, a targeted SCADA attack requires process knowledge that the outside hacker or malware author usually doesn't possess.

4.3 Target Qualification Beyond Critical Infrastructure

Most discussions on SCADA attacks have focused on critical infrastructure. It can be assumed that this is primarily influenced by feared terrorist attacks with the aim to disrupt civil life. While the cyber terrorist threat was discussed in the preceding section, it seems clear that an attacker with no terrorist background would benefit little from attacking critical infrastructure. There are, however, targets in industries not belonging to critical infrastructure that seem to be at risk. Mainly these are commercial production facilities with fully or highly automated production and products or procedures that are valuable for the attacker – either in destructive or in non-destructive scenarios. For example, suppliers in the automotive industry with its just-in-time/just-in-sequence supply chains constitute what the military calls a "target rich environment". The same is true for highly automated plants in the food & beverage and pharmaceutical industries. It looks like those potential targets have been overlooked by the discussion on SCADA security in the past, even though attack scenarios exist that may even extend damage beyond the single corporation under attack.

5 Conclusion

Much of the past research on SCADA security has focused on vulnerabilities. While excellent work in this area has been done, it has to be acknowledged that so far, the overall increase in SCADA security as seen on the plant floor is small. This is partly due to the fact that while SCADA vulnerabilities are very clear to most, threats are not. If we do assume that the biggest threat to SCADA security is intentional attacks, the authors suggest to explore and investigate this threat using a scenario-based approach, as we cannot base our investigation on empiric evidence. A scenario-based approach has several benefits:

- With an idea of what SCADA attacks would actually look like, we can establish the likelihood of such attacks for a given target by matching the target's environment with the attack scenario's characteristics.
- If an attack looks like it would succeed for a given target, we may use the attack scenario's characteristics to furnish a specific counterstrategy.

- Last but not least, we may use the attack scenarios to make risk plausible to decision makers and stakeholders, which is crucial for obtaining the funding and support required for risk management.

Some scenarios presented in this paper did happen, others will happen. The authors nevertheless assume that SCADA attacks are unlikely for most organizations. For those that may be at risk, scenario-specific countermeasures have been suggested that can help a great deal of reducing such risk. It should, however, be noted that there are other, non-intentional threats to SCADA installations, such as the random, SCADA-unaware, malware threat or the threat of accidental misconfigurations. While not the subject of this paper, non-intentional threats account for the majority of real-life incidents and materialize into damage on a regular basis. Even if non-intentional threats are not quite as thrilling and dramatic as malicious attacks, it is suggested that they are given the same attention. Non-intentional threats can and should also be presented as scenarios, if just to gain better acceptance from decision makers.

It could be argued that the scenarios presented here may serve as an inspiration for real-world attackers. The authors believe that such attackers will figure out successful attack strategies anyway, or have already done so. The more important task for risk management is to know what to expect, just to be able to prepare in time, with a clear picture of what to protect against. With no idea about what to protect against, the value of countermeasures is stochastic. Besides that, most organizations don't even start the risk management process as decision makers simply don't see any credible SCADA threat. Nobody should expect significant efforts to reduce a risk that didn't materialize in the past and that doesn't even sound plausible. Scenarios can at least help to establish plausibility, which often is a crucial point to establish SCADA security as a business case.

The goal of this paper was not to detail every possible risk scenario. The goal was to lay the fundaments of a SCADA threat modeling framework that is illustrated by sample attack scenarios. The authors hope that other researchers, and especially the people responsible for risk management in potential target organizations, may use this framework to determine other risk scenarios that asset owners should protect against, and to suggest which countermeasures to apply best.

About the Authors – Ralph Langner has accumulated 20 years experience in industrial IT and networks. He is founder and CEO of Langner Communications AG, a software and consulting company in Hamburg, Germany. Ralph has a broad experience with risk assessments and security consulting assignments for corporations in the manufacturing and process industries. He has published several articles and research reports on SCADA security and is a member of the VDI/VDE standards body that develops the upcoming German security standard 2182. For more information on Langner Communications, please refer to www.langner.com.

Bryan Singer is Vice President of Professional Services with Wurldtech Security Technologies (http://www.wurldtech.com) and Co-chairman of ISA-99 Industrial Automation and Control Systems Security standard. He began his professional career with the US Military focusing on

issues such as physical, systems, network security and force protection. Since that time, he has worked in software development in over 25 professional coding languages, worked in UNIX and mainframe systems, supported large scale ERP, MES, LIMS, and SPC implementations, and has spent significant time in cyber security projects focusing on risk analysis, vulnerability testing, penetration testing, risk mitigation strategies, and enterprise architecture and design including technical and policy based countermeasures and remediation strategies. Mr. Singer holds the CISSP and CISM certifications, and runs a personal blog site at http://www.cipiq.com and is an active contributor to http://www.wurldtech.com/blog.

References

[1] Roper, Carl A., Risk Management for Security Professionals, Butterworth-Heinemann, 1999.

[2] Schneier, Bruce, Secrets and Lies: Digital Security in a Networked World, Wiley, 2004.

[3] Swiderski, Frank and Snyder, Window, Threat Modeling, Microsoft Press, 2004.

[4] Schneier, Bruce, "Modeling Security Threats", Dr. Dobbs Journal, December 1999, http://www.schneier.com/paper-attacktrees-ddj-ft.html.

[5] Kaczynski, Theodore, "Industrial Society and Its Future", http://en.wikisource.org/wiki/Industrial_Society_and_Its_Future.

[6] US Army Training and Doctrine Command, "Cyber Operations and Cyber Terrorism", DCSINT Handbook No. 1.02, August 2005 http://www.hitechcj.com/sitebuildercontent/sitebuilderfiles/us.army.guide.supp.two.pdf

[7] Singer, Bryan, "Interesting Fun on ISC", CIPIQ blog, November 2007 http://www.cipiq.com/2007/11/09/interesting-fun-on-isc/ .