INL/CON-07-12849
PREPRINT

Ideal Based Cyber Security Technical Metrics for Control Systems

2nd International Workshop on Critical Information Infrastructures Security

Wayne Boyer
Miles McQueen

October 2007

The INL is a
U.S. Department of Energy
National Laboratory
operated by
Battelle Energy Alliance

Ideal Based Cyber Security Technical Metrics for Control Systems

Wayne Boyer[1], Miles McQueen[1]

[1] U. S. Department of Homeland Security, Control System Security Center,
Idaho National Laboratory, 2525 Fremont Ave.,
Idaho Falls, Idaho, USA 83404
{wayne.boyer, miles.mcqueen}@inl.gov

Abstract. Much of the world's critical infrastructure is at risk from attack through electronic networks connected to control systems. Security metrics are important because they provide the basis for management decisions that affect the protection of the infrastructure. A cyber security technical metric is the security relevant output from an explicit mathematical model that makes use of objective measurements of a technical object. A specific set of technical security metrics are proposed for use by the operators of control systems. Our proposed metrics are based on seven security ideals associated with seven corresponding abstract dimensions of security. We have defined at least one metric for each of the seven ideals. Each metric is a measure of how nearly the associated ideal has been achieved. These seven ideals provide a useful structure for further metrics development. A case study shows how the proposed metrics can be applied to an operational control system.

Keywords: Cyber Security Metrics, Control System Security.

1 Introduction

Electronic control systems are used to operate much of the world's critical infrastructure and are increasingly connected to public networks. Therefore, control systems and the associated critical infrastructure are at risk from cyber attacks. Examples of critical infrastructures that may be at risk from cyber attack are power plants, chemical processing plants, rail and air transportation, oil and gas facilities, etc. The security of a control system (or of any electronic network) is difficult to quantify. Meaningful metrics are needed to make informed decisions that affect system security.

A metric is a standard of measurement. The goal of metrics is to quantify data to facilitate insight [5]. It is important which metrics are chosen because good metrics lead to good decisions and bad metrics lead to bad decisions. The scope of this paper is limited to quantitative technical metrics. A cyber security technical metric is the security relevant output from an explicit mathematical model that makes use of objective measurements of a technical object. Other types of metrics (such as operational and organizational metrics and metrics that are qualitative such as "low

impact" or "highly unlikely") can provide insights about security but are beyond the scope of this work.

The overarching goal of technical metrics is the estimation of risk where risk is defined as the probability of an event times the consequence of the event. Security risk is generally stated as equal to the Threat times Vulnerability times Consequence. The risk we would like to measure is the expected value of the loss from cyber attacks per unit time. The estimation of risk could provide the ability to weigh the benefits versus costs of security counter measures.

Previous work [6] proposed "mean time-to-compromise" as a security metric and proposed a simple method for calculating it as a function of the number of known vulnerabilities. A method was also proposed for estimating risk reduction for a simple control system using the mean time-to-compromise metric [7]. However, those methods require simplifying assumptions that are not valid in general. A credible estimation of cyber security risk in real world control systems is not currently feasible because the problem involves an unpredictable intelligent adversary and very complex systems. The metrics we propose in this paper are intended to support the concept of risk measurement within the practical constraints of what is currently objectively measurable and what is potentially under the control of the defender. A good set of metrics should have the following attributes: The number of metrics should be small (less than 20) to be manageable; the metrics should be easy to understand, measurable and objective; the metrics should be directly related to security risk; and the set of metrics should represent the most important measurable security attributes of the system.

2 Initial security metrics investigation

Thirty guides and standards documents (including, for example, references [2], [3], [12], [13]) were reviewed in search of technical metrics that have previously been defined and recommended [4]. A sampling of security metrics used by some industries were also included in the investigation. Most of the metrics found in the standards and guides do not meet our definition of a technical metric. We found no case where a standards document recommended the use of a specific metric or set of metrics.

We evaluated the strengths and weaknesses of the few identified technical metrics and concluded that existing metrics have serious weaknesses. For example, many of the metrics were simply a percent of the system components that implemented a certain type of security control mechanism. But the fractional implementation of a given security mechanism does not necessarily correlate to risk. A specific metric defined in industry is "Average number of vulnerabilities per system component". This metric has the following strengths. It is easy to understand and it easy to obtain estimates by automatic scanning tools. But the problem of using an average is that all vulnerabilities and all components of the network are given equal weight. Consider the case where there is one easily exploitable vulnerability that allows penetration of a critical system component while there are zero known vulnerabilities on the other system components. Now consider a case where there are no known vulnerabilities on

critical components, no vulnerabilities that allow penetration from an external site, but there are many minor vulnerabilities on non-critical system components. The former case is a high-risk situation, but the metric indicates low risk while the latter case is a low-risk situation, but the metric indicates high risk. This metric has a built in assumption that all vulnerabilities and all components are of approximately equal value. The assumption is false for most systems. The metric can be improved by averaging the number of vulnerabilities for each group of components with similar security implications and for vulnerabilities with similar effects (i.e., external penetration versus privilege escalation). The results of our investigation of existing metrics showed the need for the definition of a small set of technical metrics that operators of control systems can use to gain better insight into their security risk.

3 Approach

The measurement of risk is the overarching goal of security metrics but is currently highly subjective. Since a credible estimate of risk is not feasible, we suggest a set of ideals to guide the development of a set of objective measurements that can provide decision makers with improved insights about security risk.

Table 1. Seven abstract dimensions of security and associated ideals

Security Dimension	**Ideal**
1. Security Group (SG) knowledge	1. Security Group (SG) knows current system perfectly.
2. Attack Group (AG) knowledge	2. Attack Group (AG) knows nothing about the system.
3. Access	3. System is inaccessible to AGs
4. Vulnerabilities	4. The system has no vulnerabilities
5. Damage potential	5. The system can't be damaged
6. Detection	6. SG detects any compromise instantly.
7. Recovery	7. SG can restore system integrity instantly.

3.1. Seven ideals of security

Seven ideals are the basis for our proposed metrics. Each ideal is associated with an abstract dimension of security and represents a system condition at a given point in time such that perfection has been achieved for its associated dimension of security. The seven dimensions of security and the respective ideals are listed in Table 1. We chose the ideals in Table 1 based on our study and experience in the cyber security field and suggest that each of these ideals is strongly related to security risk.

3.2. Security principles

It is well known that the purpose of computer security is the protection of confidentiality, availability and integrity of computer systems. Security principles support that purpose. We assert that our seven security ideals are consistent with generally accepted security principles. To support that assertion we successfully mapped security principles from Bishop [1], Neumann [10], Schneier [14], NIST [16] and Summers [17] to our seven ideals.

To help identify a useful set of technical metrics we suggest the following set of principles that are organized by and directly applicable to our seven abstract dimensions of security.

1. Security Group (SG) knowledge principles
   a. The system configuration should not be changed without the security group's knowledge.
   b. The system should be thoroughly tested and regularly monitored for vulnerabilities.
2. Attack Group knowledge principles
   a. Credential keys (e.g. passwords) should be strong, should not be disclosed and should be changed regularly.
   b. The system should send no unencrypted information through external networks or respond to any user/application/machine that has not previously been authenticated.
   c. Information about the system design, implementation or configuration should not be made public.
3. Access principles
   a. Number of external communication paths should be minimized; including network connections, TCP/IP ports/services, physical access to USB ports and portable storage media drives.
   b. Compartmentalization. The system should be divided into loosely coupled parts. This principle improves security because if one part is compromised, the damage to the rest of the system is limited. This principle avoids total loss from a single point of failure. The principle includes the precept of least privilege.
   c. Defense in depth. The system should be designed and configured such that an attack can succeed only by breaking through a series of independent barriers.
4. Vulnerability principles
   a. The time between vulnerability discovery and repair should be small.
   b. Complexity implies unknown vulnerabilities.
   c. Fix high-priority vulnerabilities first, with priority on vulnerabilities that can be exploited from the perimeter and that allow penetration.
5. Damage potential principle
   a. Mechanisms that are independent of the control system should provide protection such that the cost of damage due to control system malicious behavior is minimized.
6. Detection principles

a. The system should be constantly monitored for malicious behavior and alarms should be raised when detected.
b. The malicious behavior detection mechanisms must not have false positive rates that exceed the ability of the SG to process, even if this results in some malicious behaviors going undetected.

7. Recovery principles
a. Several previous versions of system data should be saved regularly and protected from deliberate or accidental loss, such that in the event of compromise, a previous version can be chosen that is not likely to be corrupted.
b. The time needed to restore the system with a previous uncorrupted version should be small.

4. Proposed set of metrics

Table 2. Proposed metrics

Security Ideal	Metric	Principle
1. Security Group (SG) knows current system perfectly.	Rogue change days	1a
	Component test count	1b
2. Attack Group (AG) knows nothing about the system.	Minimum password strength	2a
	Data transmission exposure	2b
3. System is inaccessible to AGs	Reachability count	3a
	Root privilege count	3b
	Defense depth	3c
4. The system has no vulnerabilities	Vulnerability exposure	4a
	Attack surface	4b
5. The system can't be damaged	Worst case loss	5a
6. SG detects any compromise instantly.	Detection mechanism deficiency count	6a
	Detection performance	6b
7. SG can restore system integrity instantly.	Restoration time	7b

Our proposed metrics are based on the seven security ideals listed in Table 1. We propose at least one metric for each of the seven ideals as shown in Table 2. Each defined metric is intended to answer the question "what can be objectively measured on the system that is a reasonable representation of how nearly the ideal has been realized?" The following sections briefly discuss each of our proposed metrics.

Rogue Change Days is the number of rogue changes multiplied by number of days the changes were unknown to the Security Group (SG). A rogue change is any change to the system configuration without prior notification to the SG.

A key assertion is that the security risk from changes to the system without notification to the security

group is, on average, worse than for changes which are announced in a well managed system.

This metric is a valid worst case measure of the quantity of potentially security impacting changes. One weakness of this metric is that it does not include any measure of the actual security impact of changes.

For this metric the set of objects under change control must first be established and a version identifier must be saved for each object to establish a baseline. Periodically the current version identifier is scanned and compared to the previously saved identifier. Examples of objects under configuration management are: PLCs, HMIs, critical computer files, network devices attached to the local network, etc.

Each type of configured object must have an associated mechanism for identification that produces an identifier that an audit program can obtain from the object. For example, computer files may have a hash function applied to the file content to calculate an identifier that can be used to determine if the file has changed.

Mathematical formula:

S_T == An ordered set of version identifiers for all configured objects, measured at time T.

S_{T+k} == An ordered set of version identifiers for all configured objects, measured at time T + k.

TSC_{T+k} == Number of mismatches between sets S_T and S_{T+k}

CC_{T+k} == Changes introduced into the system only after notification of the security group,

RC_{T+k} == TSC_{T+k} - CC_{T+k} is the number of Rogue Changes between the current measurement of the system and the previous measurement of the system.

Rogue Change Days == $RC_{T+k} * k$

Component Test Count is the number of control system components that have not undergone independent security testing. This metric is included in our proposed set because we recognize the importance of security testing. A key assertion is that independent security testing of the system components will reduce risk.

An independent test is one that is performed by personnel that are not under the direct employ of the vendor. An unresolved question: Do tests become obsolete with the passage of time or when there is a new version of the component? If so, then how do you determine when the tests are obsolete?

Minimum password strength is the shortest time (in days) needed to crack a single password for any account on the system.

Key assertions are that passwords are the most critical information to protect on the system and the system security tends to improve when minimum password strength increases. This metric is a valid measure of the minimum amount of time an attacker would need to compromise the system by password cracking. The password age should be subtracted from the password cracking time. One weakness of this metric is that it does not measure the strength of other authentication mechanisms but passwords are the most common form of authentication.

Data collected for this metric is the encrypted password files from all machines on the system. For example, all password files from UNIX servers, Configuration data

for Web Servers, Database Servers, Windows workstations, Control System HMI, etc. A password cracking tool is then applied to each password file instance. The metric is simply the minimum time needed to crack a single password.

Password cracking tools are available commercially and for free download. Data should be collected whenever passwords change. This metric is an important measure because passwords (digital private keys) are by far the most common form of authentication. The value of the metric should be greater than the password expiration time. This metric is independent of password policies because it measures the least amount of time an attacker would need to crack a password if the encrypted password data is available to the attacker. If a very weak password is used, (including a default vendor supplied password) an attacker can guess the password without obtaining the encrypted password files and this metric would detect that high risk situation because good password cracking tools crack very weak passwords virtually instantly. Passwords used for authentication at the perimeter are particularly important and therefore perhaps should be measured for strength separately from other passwords used on the system. The security manager should ensure that vendor supplied passwords and passwords commonly used by maintenance personnel are included in the password cracker's dictionary.

Data transmission exposure is the unencrypted data transmission volume. A key assertion is that any data that can be monitored by a potential attacker increases the security risk. Some data is clearly more sensitive than others but to make the metric simple to obtain we propose that this metric be a count of the number of unencrypted machine communication channel pairs in use. For a TCP/IP network, it is the number of unencrypted machine TCP-port pairs in use (as observable by network monitoring). Some network paths are more critical than others but during a multi-stage attack, an attacker may gain access to an internal network by first penetrating the system through an external network path. The security manager may choose to categorize network connections (e.g. publicly accessible, internal) and track this metric for each network category.

Reachability count is the number of access points (relative to a specific point of origin such as the Internet). A key assertion is that a reduction in the number of access points tends to reduce the cyber security risk.

This metric requires complete network configuration information including connectivity and firewall rules. It also requires information about physical access to computer ports. The system may be scanned to identify all network communication paths. Physical access to portable storage media drives can be done by inspection.

Mathematical model:

N_s == Number of ports (services) that respond to data transmitted from the point of origin.

N_o == Number of machines that have network connectivity from inside the network to the point of origin. Connectivity means the network configuration allows the machine to originate two-way connection-oriented sessions to some facility located at the point of origin. (Note: strict one-way outgoing data transmission is OK)

N_p == Number of physical access points to unrestricted portable storage media drives.
N_T == Total reachability count

$N_T = N_s + No + Np$

The security manager may choose to combine the network and physical reachability counts or track them separately.

Because of the possibility of penetration of the perimeter the security manager may choose to calculate this metric at multiple points of origin within the network perimeter such as at the DMZ, or behind each firewall. The measurement of reachable ports/services includes all the cases of crafted packets that exploit known vulnerabilities in firewalls and routers, such as the spoofing of IP addresses and packet fragmentation to disguise the targeted TCP port number.

The point of origin for physical access may be "outside the fence" or some other partially controlled area or combination of areas within the fence as defined by the security manager. Examples of restricted portable storage media drives that should not be included in the count of physical access points are:

- USB ports that are disconnected or physically disabled.
- Host-based or device-based port encryption.
- Ports restricted by end-point control software.

Root privilege count is the number of unique user IDs with administration (root) access privilege. A key assertion is that risk is strongly related to the principle of least privilege. This metric is a simple measure of how well this principle is being followed.

Defense depth the minimum number of independent single machine compromises required for a successful network attack. This metric emphasizes the need to avoid a protection configuration that can be defeated by a single point of failure. There may be common vulnerabilities on various paths of entry, therefore the attack steps may not be truly independent and this metric may be optimistic. To calculate this metric detailed network configuration data is needed such that each machine in the system can be determined to be reachable or not reachable from every other machine and every network access point in the system. A machine is defined to be reachable from a point of origin if at least one service responds to data transmitted from that point.

Mathematical model:

Defense Depth == Minimum number of compromises required to reach any machine in the set S from the public network by traversing network paths. S is the set of machines such that if any machine in the set is compromised then the attack is considered to be successful.

Vulnerability exposure is the sum of known and unpatched vulnerabilities, each multiplied by their exposure time interval. It is measured in vulnerability days. A key assertion is that the longer a vulnerability is open the greater the risk it will be exploited.

Mathematical model:

N = Number of open known vulnerabilities that apply to the system.
T_i = Discovery date of vulnerability i
t = current date
T == Total vulnerability days

$$T = \sum_{i=1}^{N} (t - Ti)$$

For publicly disclosed vulnerabilities, the discovery date is the disclosure date from the public vulnerability database. For vulnerabilities that are discovered locally, such as configuration errors, it is the local discovery date. Vulnerabilities that apply to the system may be identified by vulnerability test tools and by comparing system components to the components associated with publicly disclosed vulnerabilities.

The system should be scanned for vulnerabilities often (suggest weekly or when there is a known configuration change). Public vulnerability databases should be checked regularly and often (suggest daily). This metric is affected by vulnerability discovery rate and by patch rate. Vulnerabilities may result from design errors, implementation errors and from mis-configurations such as inappropriate trusted relationships between machines. Some vulnerabilities are more significant than others. Tools such as Attack Graphs [11] can be used to determine priority categories for all known vulnerabilities. The Common Vulnerability Scoring System (CVSS) [15] is another suggested mechanism for prioritizing known vulnerabilities. This metric could be applied separately for each vulnerability category.

Attack surface is a measure of potential vulnerability. Key assertions are 1) vulnerabilities exist that are currently unknown to the defender and 2) the attack surface complexity, including external interfaces is strongly correlated to the potential for the discovery of new vulnerabilities.

Attack surface has been proposed as a security metric by Manadhata and Wing [9]. This metric is considered to be potentially very valuable but is not yet sufficiently developed to be used in practice.

Worst case loss is the maximum dollar value of the damage/loss that could be inflicted by malicious personnel via a compromised control system.

A key assertion is that the risk is strongly related to the worst case loss. Although there can be successful attacks where the actual loss is much less than worst case, we assert that a reduction in the worst case loss reduces the potential for loss and therefore reduces risk. The worst case loss can probably be estimated from an existing safety analysis associated with the plant. The metric is the answer to the question "If the control system is under the control of a malicious person, what damage can be done?". Safety systems that prevent serious damage should be completely independent of the control system.

Detection mechanism deficiency count is the number of externally accessible devices without any malware/attack detection mechanisms. A key assertion is that detection mechanisms reduce risk especially when applied to devices that can be used as entry points for attacks.

Detection performance is a measure of the effectiveness of the detection mechanisms (intrusion detection system, anti-virus software, etc.) implemented on the system. The metric can be defined as detection probability discounted by false alarm rate.

The metric should be applied separately to each of the detection mechanisms used on the system.

A suggested mathematical model:

N = Number of attack test cases

D = Number of attack test cases detected

Pd = D/N = Probability of detection.

F = Number of false alarms during tests.

Pfa = F/(D + F) = Probability of false alarm.

Detection Performance = Pd * (1 – Pfa)

This metric is difficult to obtain currently but is theoretically measurable. There is some public data available but better tests and tools are needed. Some intrusion detection products have been evaluated by Lincoln Laboratories [8].

Restoration time is the worst case elapsed time to restore the system to a known uncorrupted version. The metric can be determined by measuring the actual time elapsed from "start" to a fully restored and 100% operational system. If it is impractical to perform that kind of a test on an operational system then this data should be collected for actual security events if they have occurred. If a recovery test is not feasible, then a worst case recovery analysis may be used to estimate recovery time.

T_0 = Start time (Time compromise is detected, or test start time)

T_r = Time at which recovery is complete and the system is 100% operational.

Restoration time = Maximum value of all instances of (T_r -T_0)

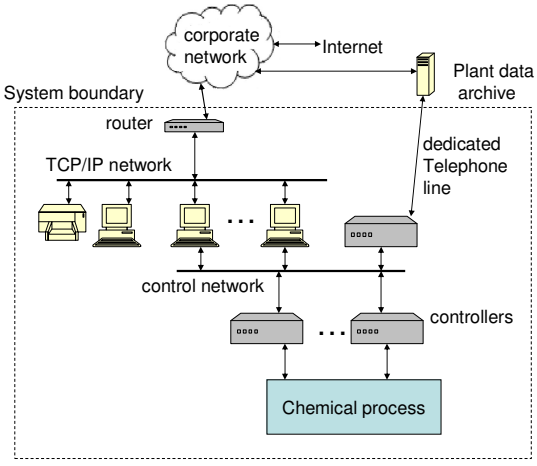

Fig. 1. Case study control system network diagram

5. Case Study

Our proposed security metrics were applied to a case study of a Distributed Control System (DCS) for a chemical processing plant. Figure 1 is a simplified network diagram of the case study system. Notice that the system is connected to the Internet through the corporate network. The

router that provides connectivity between the corporate network and the local TCP/IP network restricts access to the control system with an access-control-list so only the incoming TCP/IP connections with origination addresses that match the control list are allowed through the router. The system boundary was defined to be the processing plant and the control system networks that are within the control room. The corporate network affects the security of the control system but for this study the corporate network was not considered to be part of the system.

Table 3. Case study metrics values

Metric Name	Metric Value	Ideal target value	Suggested target value
Rogue Change Days	0	0	0
Minimum Password Strength	> 30 days	∞	>30 days
Data Transmission Exposure	23	0	1
Reachability Count (N_T)*	164	0	1
Physical (N_p)	2	0	0
Services (N_S)	149	0	1
Outgoing (N_o)	13	0	0
Root Privilege Count	3	0	1
Defense Depth	2	∞	4
Worst Case Loss	$100M	$0	?
Detection Mechanism Deficiency Count	12	0	0
Vulnerability Exposure (high priority)	16,416 vuln. days	0	0
Vulnerability Exposure (low priority)	15,877 vuln. days	0	0
Restoration Time	120 minutes	0	120 minutes

*$N_T = N_p + N_S + N_o$

The DCS for this case study consists of a TCP/IP network that provides connectivity for 11 workstations and 2 printers, and a proprietary control network that provides connectivity to approximately 30 distributed controller nodes that control and monitor the plant. The workstations on the TCP/IP network consist of standard IT hardware, standard IT operating system software and application software supplied by the DCS vendor. The controller nodes consist of specialized control hardware and software supplied by the same DCS vendor.

5.1. Metrics not included in the case study

The values of the following proposed metrics were not obtained for the case study. Not surprisingly, the attempt to determine the values of these metrics showed that these metrics are difficult to measure. These metrics are currently impractical, but remain in the proposed set because they are theoretically measurable and may become practical in the future as more advanced tools are developed.

- Component Test Count
- Attack Surface
- Detection Performance.

5.2 Case study metrics values

The metric values, ideal target values and suggested target values for our case study are shown in Table 3. The "suggested target value" was determined by estimating what the value of the metric would be after making a set of suggested security improvements. The cost of the suggested improvements can now be weighed against the value the projected improvements in the metrics. Every suggested security improvement will result in the improvement of at least one of the recommended metrics. The method for obtaining each metric value and suggested security improvements are described below.

5.2.1. Rogue Change Days The case study system has an audit mechanism that compares the system configuration to the official database of configured items. There have been no known cases of a rogue change on this control system. Therefore, the measured value for the metric is zero.

The system has a configuration management plan that has identified a long list of configured items of many different types including all hardware and software items related to the options that apply to this system, such as the set of display screens, Control-Language Programs, Tags and history parameters. The audit program resides on a workstation that is located outside the system boundary. The system administrator runs the audit program after system configuration changes are made to verify that only the planned changes have taken place. The audit program could be fooled by a clever attacker because it primarily compares file dates to the list of configured item file dates. Additional tools could be used to provide more reliable measures of whether there have been unauthorized changes to the system.

5.2.2. Minimum password strength The system did not use any default passwords. The age of all the passwords was 2 days. (passwords were all changed 2 days before the case study started). The password files were copied from all workstations on the TCP/IP network and a freeware password cracker (John the Ripper) was run against the password files. The password cracker ran for 30 days without cracking any passwords, therefore the value of the metric is greater than 30 days. Since the system administrator sets all passwords and uses a password policy that includes a minimum number of characters, the passwords for this system seems to be quite strong.

5.2.3. Data transmission exposure The monitoring of network traffic at the router on the system boundary showed that several unencrypted services are used including DNS, remote login, print services and FTP. There are 11 machines on the local TCP/IP network that use the DNS service located outside the control room, 9 machines on the TCP/IP network provide remote login and FTP services, there are 2 printers on the TCP/IP network that provide print services to external hosts. The dedicated telephone line that provides data to the plant data archive was counted as one data transmission machine-port pair. The total number of machine-port pairs was 23. This metric could be reduced significantly by setting up a firewall that allows no unencrypted traffic from the TCP/IP network to the corporate network. Needed services could be provided by proxy servers and encrypted services. The suggested

target value of 1 reflects the fact that it may not be feasible to encrypt the data that flows to the plant data archive.

5.2.4. Reachability count The network services reachability count was obtained by scanning the machines connected to the local TCP/IP network with the well known open source tool Nessus. Each unique machine type was scanned and then the total numbers were obtained by adding the number of reachable services on every machine of each type. This metric can be reduced by turning off unneeded services however it may be difficult to determine which services are not needed. A firewall at the control room boundary that allows only secure shell service to be accessed externally would allow this metric to be reduced to the value of one and would clearly improve the security. We suggest that all needed externally accessible services could be provided through the secure shell service by some changes in system configuration.

The outgoing reachability count was obtained by simply counting the number of machines on the local TCP/IP network because there are currently no outgoing restrictions. We suggest that all machines should be restricted by a local firewall to disallow all outgoing connections. This restriction would change the metric to a value of zero and would clearly reduce risk from attacks that use outgoing connections such as access to external web sites as a pathway in.

The physical reachability count is the number of workstations in the control room with unrestricted USB ports. Although the control room has 24 hour per day monitoring malware could be easily introduced into the control system through an unrestricted USB port by an unsuspecting innocent user through a thumb drive. We suggest restricting the USB ports which would reduce the metric to a value of zero.

5.2.5. Root privilege count The number of unique user ID's with administrative access privilege was small (3), so this metric indicates no serious contribution to risk.

5.2.6. Defense depth Although the corporate network is outside the system boundary it affects the value of the defense depth metric because it separates the control system from the public network. The minimum number of stages for a successful attack is two for our case study under the assumption that an attacker must first gain access to the corporate network and then compromise one of the machines on the local TCP/IP network. The engineering workstation and operator consoles are connected to the TCP/IP network, therefore a compromise of any of those machines would constitute a successful attack. We suggest that security would be improved and the metric value would be increased from 2 to 4 by the following actions.

Standard security practices on corporate networks include firewalls and DMZ that create some network partitioning. If these practices are followed on the corporate network then the number of stages required for an attacker to reach the control room boundary will be at least 2 which would increase the metric by one. The value of the metric could be incremented again by making the following changes in the control room: The control room network could be partitioned behind a local firewall such that an attacker could not reach any of the critical machines directly through the TCP/IP network. If the communication path from the control system to the data archive were configured to allow only one-way outgoing data transmission, that path would be removed as a possible path of attack.

5.2.7. Vulnerability exposure All the unique machine types on the TCP/IP network were scanned for vulnerabilities by the Nessus tool. There are no known tools available that scan for vulnerabilities on the control network. The vulnerability scanner identified some low priority vulnerabilities that are in the public CVE database so the discovery times for those vulnerabilities were obtained from the CVE database. Some other vulnerabilities had previously been identified on the case study system but had not been publicly disclosed so the discovery times for those vulnerabilities were obtained from the date on the memorandum that described the vulnerabilities. The vulnerabilities were categorized as either high or low priority and the metric was calculated for each category. High priority vulnerabilities allow an external penetration while low priority vulnerabilities do not. If the same vulnerability was found on more than one machine, it was counted separately for each machine. Table 3 shows that the number of vulnerability days is a large number for both vulnerability categories. This metric clearly shows the need for action. The known vulnerabilities have known mitigation methods which would improve system security and reduce the metric value to the ideal of zero.

5.2.8. Worst case loss The worst case loss for our study was estimated by the plant designers to be about $100M based on the costs of reconstruction, repair and lost production for the most extreme case of malicious behavior by the control system. This is significant and implies the need for some independent safety mechanisms.

5.2.9. Detection mechanism deficiency count The machines that qualify as "externally accessible" are all the machines that have a data transmission path directly connected to the network located outside the control system boundary. There were 13 machines connected directly to the router which connects to the corporate network. The connection to the plant data archive is also externally accessible. Therefore, the number of externally accessible machines is 14. Only 2 of the 14 machines have any malware detection. Therefore the value of the metric is 12. The value of the metric can be improved by reducing the number of directly accessible machines as suggested for improving the defense depth metric above, or by installing more detection mechanisms. For our case study, an ideal value of zero is achievable.

5.2.10. Restoration Time The restoration time for our case study system has been measured during normal preventive maintenance activities. Reboot time for the entire system was measured at the time of new software installation to be 120 minutes. This time is limited by the system architecture.

6. Conclusions

Because of the complexity of networked control systems and the unpredictable nature of intelligent adversaries, a credible quantitative measure of security risk is not currently feasible. However, the seven security ideals provide a useful structure for thinking about security and for further development of technical security metrics. A

well chosen set of metrics can help the security managers make better decisions that will lead to real security improvements. The specific metrics proposed here provide a small and manageable set that may be refined and expanded while they encourage management decisions that tend to reduce the risk of a successful cyber attack on control systems. The definition of the proposed metrics has identified the need for improved measurement tools. A case study that applied many of the proposed metrics to a real control system showed that recommended security improvements correspond to improvements in the values of one or more of the proposed metrics.

References

1. Bishop, M., Computer Security Art and Science, Addison Wesley, pp. 343-349, 2003.
2. Chew, E., Clay, A., Hash, J., Bartol, N., Brown, A., Guide for Developing Performance Metrics for Information Security, NIST Special Publication 800-80, May 2006.
3. Chemical Sector Cyber Security Program (CSCSP), Guidance for Addressing Cyber Security in the Chemical Industry, Technical Report, CSCSP, May 2006.
4. Idaho National Laboratory Report to the Department of Homeland Security, INL/EXT-06-12016, Cyber Security Metrics, December 2006.
5. Jacquith, A., Security Metrics, Addison Wesley, 2007.
6. McQueen, M. A., W. F. Boyer, M. A. Flynn, G. A. Beitel, "Time-to-compromise Model for Cyber Risk Reduction Estimation", First Workshop on Quality of Protection, Sept. 2005.
7. McQueen, M. A., W. F. Boyer, M. A. Flynn, G. A. Beitel, "Quantitative Cyber Risk Reduction Estimation Methodology for a Small SCADA Control System", Proceedings of the 39th Hawaii International Conference on System Sciences, pp. 226, Jan. 2006.
8. Mell P, V Hu, R Lippmann, J Haines, and M Zissman, An Overview of Issues in Testing Intrusion Detection Systems, Interagency Report (IR) 7007, National Institute of Standards and Technology, Gaithersburg, Maryland, June 2003,
9. Manadhata, P., Wing, J. M., An Attack Surface Metric, Technical Report CMU-CS-05-155, July 2005
10. Neumann, P. G., Computer Related Risks, Addison Wesley, pp. 244, 1995.
11. Ou, X., Boyer, W., McQueen, M., A Scalable approach to Attack Graph Generation, 13th ACM Conference on Computer and Communications Security, CCS'06, October 30 through November 3, 2006.
12. Ross, R., S. Katzke, A. Johnson, M. Swanson & G. Rogers, System Questionnaire with NIST SP 800-53: Recommended Security Controls for Federal Information Systems, Technical Report, NIST, References and Associated Security Control Mappings, Gaithersburg, Maryland, March 2006,
13. Swanson, M., N. Bartol, J. Sabato, J. Hash & L. Graffo, NIST Special Publication 800-55: Security Metrics Guide for Information Technology Systems, Technical Report, National Institute of Standards and Technology (NIST), Gaithersburg, Maryland, July 2003
14. Schneier, B., Secrets & Lies, Wiley, pp. 367-380, 2000.
15. Schiffman, M., A Complete Guide to the Common Vulnerability Scoring System (CVSS), Technical Report, Forum for Incident Response and Security Teams (FIRST), June 7, 2005.
16. Swanson, M., B. Guttman, "Generally Accepted Principles and Practices for Securing Information Technology Systems", NIST 800-14, September 1996.
17. Summers, R. C., Secure Computing Threats and Safeguards, McGraw Hill pp. 251-252, 1997.