Primer Control Systems Cyber Security Framework and Technical Metrics

June 2009

Homeland Security

Control Systems Security Program

National Cyber Security Division

EXECUTIVE SUMMARY

The Department of Homeland Security Control Systems Security Program has developed a control systems cyber security framework and a set of technical metrics to aid owner-operators in managing control systems cyber security. The framework defines seven relevant cyber security dimensions and provides the foundation for managing control systems cyber security. Based on the developed control systems cyber security framework, a set of ten technical metrics are provided that allow control systems owner-operators to track improvements or degradations in their individual control systems security posture.

This primer has been developed to assist owners and operators in securing control systems by applying the control systems cyber security framework and implementing ten technical security metrics.

CONTENTS

EXECUTIVE SUMMARY iii

1. INTRODUCTION 1

2. CONTROL SYSTEMS CYBER SECURITY FRAMEWORK 2
   - 2.1 Security Group Knowledge 3
   - 2.2 Attack Group Knowledge 3
   - 2.3 Access 3
   - 2.4 Vulnerabilities 3
   - 2.5 Damage Potential 4
   - 2.6 Detection 4
   - 2.7 Recovery 4
   - 2.8 Cyber Security Dimensions Drive Cyber Security Metrics 4

3. RECOMMENDED TECHNICAL METRICS 6
   - 3.1 Rogue Change Days 6
   - 3.2 Security Evaluation Deficiency Count 6
   - 3.3 Data Transmission Exposure 7
   - 3.4 Reachability Count 7
   - 3.5 Attack Path Depth 7
   - 3.6 Known Vulnerability Days 7
   - 3.7 Password Crack Time 8
   - 3.8 Worst Case Loss 8
   - 3.9 Detection Mechanism Deficiency Count 8
   - 3.10 Restoration Time 8
   - 3.11 Specifications of Recommended Technical Metrics 9

4. CASE STUDIES 20
   - 4.1 Case Study 1 20
   - 4.2 Case Study 2 21

5. SUMMARY 23

6. REFERENCES 24

FIGURES

Figure 1. Simplified network diagram of Case Study 1 (DCS). 20

Figure 2. Simplified network diagram of Case Study 2 (SCADA) 21

TABLES

Table 1. Seven Control Systems Cyber Security Dimensions. 2
Table 2. Control Systems Cyber Security Dimension Checklist and Potential Metrics. 5
Table 3. Rogue Change Days metric specification. 9
Table 4. Security Evaluation Deficiency Count metric specification. 10
Table 5. Data Transmission Exposure metric specification. 11
Table 6. Reachability Count metric specification. 12
Table 7. Attack Path Depth metric specification. 14
Table 8. Known Vulnerability Days metric specification. 15
Table 9. Password Crack Time metric specification. 16
Table 10. Worst Case Loss metric specification. 17
Table 11. Detection Mechanism Deficiency Count metric specification 18
Table 12. Restoration Time metric specification. 19
Table 13. Case Study 1 metric values. 21
Table 14. Case Study 2 metric values. 22

1. INTRODUCTION

Electronic control systems that operate much of the Nation's critical infrastructure are increasingly connected to public networks, including the Internet. Consequently, control systems and the associated critical infrastructure are at greater risk than before from externally initiated cyber attacks. The Department of Homeland Security (DHS) Control Systems Security Program (CSSP) funded development of a control systems cyber security framework to aid owner/operators in managing their control system's cyber security posture. CSSP also developed a supporting set of technical metrics intended to allow security management of factors affecting risk and mitigation of risk.

This primer describes the framework and supporting technical metrics for owners and operators of control systems environments. Rigorous application of the framework, when thinking about a control system's cyber security posture, and consistent measurement and evaluation of the metrics will support more effective cyber security decisions for control systems.

The framework consists of seven control systems cyber security dimensions. Each of these dimensions is related to risk. Based on case studies and review of control systems cyber security assessments, there is strong evidence that the framework's seven control systems cyber security dimensions capture many of the system attributes, which correlate with a control system's risk exposure. As a result, the seven control systems cyber security dimensions are recommended as a foundation for thinking about the cyber security of control systems and are also the foundation for the ten technical metrics recommended in this primer. It is these ten technical cyber security metrics that provide support for assessing whether cyber security risk exposure is increasing or decreasing over time.

A metric is a standard of measurement.[1] This primer presents ten useful and practical 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 cyber security insights, but are beyond the scope of this primer.

The remainder of this primer consists of three sections. Section 2, "Control Systems Cyber Security Framework," discusses the framework of seven control systems cyber security dimensions to aid in thinking about the cyber security of control systems; Section 3, "Recommended Metrics," presents and recommends one specific set of technical metrics that have been applied in case studies and demonstrated their value in improving control systems cyber security; and Section 4, "Case Studies," which discusses two case studies where the recommended technical metrics were applied to operational control systems.

2. CONTROL SYSTEMS CYBER SECURITY FRAMEWORK

The control systems cyber security framework consists of seven cyber security dimensions and provides the foundation for the metrics recommended later in this primer. *Each of the seven dimensions of cyber security represents an important aspect of the control system's cyber security posture at a given point in time*. There is at least one recommended metric for each dimension. An ideal value associated with each metric indicates the best that could possibly be attained for that metric. The ideal values are provided as a target to move towards, not necessarily as a target that can be fully and completely met. The seven dimensions of control systems cyber security are defined in Table 1 and then discussed more fully in the rest of this section.

Seven Dimensions of Control Systems Cyber Security:

1. Security group knowledge
2. Attack group knowledge
3. Access
4. Vulnerabilities
5. Damage potential
6. Detection
7. Recovery

Table 1. Seven Control Systems Cyber Security Dimensions.

Security Dimension	Definition
1. Security Group Knowledge	Aspects of the system or associated management processes that impact the security group's ability (i.e., the people who are directly responsible for the cyber security of the control systems) to know the system and manage changes. It includes: (a) aspects of the system and processes associated with configuration management, (b) tools (or lack of tools) that support the tracking of changes, and (c) the collection and analysis of system logs and forensics.
2. Attack Group Knowledge	Attributes of the system, processes, or actions that provide potential attackers with means to gain information about the system. It includes: (a) software defects or configuration settings that return information when the system is probed by an unauthenticated user, (b) any information about the system that can be obtained from public sources, and (c) design or implementation weaknesses that allow a user with no authenticated privilege to gain information by listening on communication paths.
3. Access	Attributes of the system design, configuration, or deployment that provide a potential attacker with the ability to send or receive data to/from a component of the control systems from the attacker's location. This includes: (a) physical access to control systems components, (b) access to control systems components through external and internal networks, and (c) access from internal components that may have been compromised. Access does not address whether or not the communication channel can be used to gain any useful information or whether sending data can provide the attacker with any desired result.
4. Vulnerabilities	Defects or weaknesses in the control systems that can be exploited to gain unauthorized privilege. This excludes defects that allow information to be obtained once access is gained without also explicitly gaining privilege. If a single defect allows an attacker to gain information and also gain privilege, that defect is defined to be a vulnerability.
5. Damage Potential	The amount of loss that a malicious attacker has the power to cause once they have gained privilege on a control system. It does not include any weaknesses associated with the process of gaining malicious control. Although actual damage may be reduced by a quick response to an attack, this dimension does not include any effects associated with attack detection or control systems recovery.
6. Detection	The ability to detect attacks and provide timely notification. This includes anti-virus software, Intrusion Detection Systems (IDS), Intrusion Prevention Systems IPS and system logging, and whether detection mechanisms are in place and the effectiveness of those mechanisms.
7. Recovery	The ability to restore control systems from a compromised state to an uncompromised state. It includes the reliability of the backup and restore facilities and the time required to recover from an attack.

2.1 Security Group Knowledge

The first control systems cyber security dimension is Security Group Knowledge. The security group represents those people in an organization who are directly responsible for the cyber security of the control systems. Security risk is strongly correlated with the security group's knowledge of control systems environments. In an ideal situation, the security group has perfect knowledge of the system, including all the hardware and software components, network topology, communication paths, normal operational behavior, and even vulnerabilities. That knowledge is needed for a security group to effectively make security decisions that protect the control systems from potential attackers. Any changes that occur to the control systems without the security group's knowledge may inadvertently introduce new vulnerabilities into the system and inhibit the introduction of reasonable mitigation measures. Perfect knowledge of the system implies a configuration management process that includes the security group in the planning of all changes and provides a mechanism for alerting the security group to any unauthorized changes.

2.2 Attack Group Knowledge

The second control systems cyber security dimension is Attack Group Knowledge. The attack group represents any of the many potential adversaries in the world who might have interest in attacking the facility through cyber means. Cyber security risk from targeted attacks is kept down when potential attackers are unable to obtain any information about the targeted control systems. Ideally, anyone who is not authorized to use a control system should be prevented from gaining knowledge of its design or configuration and be unable to obtain any information that would allow them to plan and execute an attack. This includes information an attacker might gain about a control system after they have compromised portions of it and information they might gain from other sources before the attack commences (e.g., a vendor's web site touting the targeted facility as a success story).

Be aware that even approved users may become members of an attack group when their actions on the system go beyond what they are authorized to perform, whether inadvertently or intentionally (i.e., the "insider threat").

2.3 Access

The third control systems cyber security dimension is Access. Even though authentication mechanisms are designed to prevent unauthorized use of data transfer paths, the existence of every path, authenticated or not, negatively impacts cyber security risk. The ideal situation from a cyber security perspective is to disallow any communication channels between the control systems and any location where there are potential attackers. Although achievement of this ideal is not practical in most cases, the ideal includes the absence of any electronic connections between the Internet and the control systems.

2.4 Vulnerabilities

The fourth control systems cyber security dimension is Vulnerabilities. A vulnerability is defined as any weakness or defect in the system that provides a potential attacker with a means to gain privilege intended for authorized users only. An exploit of a vulnerability leads to a compromise. An ideal system has no weaknesses and no defects. Unfortunately, all real systems have weaknesses and if an attack group is targeting the facility, they will be actively searching vulnerability disclosure sites and using techniques such as reverse engineering to find those weaknesses.

2.5 Damage Potential

The fifth control systems cyber security dimension is Damage Potential. The ideal control systems environment prevents physical damage even if electronic networks are completely compromised by an attacker. Since risk is the expected value of loss, the damage potential is directly proportional to risk. The amount of damage that can be caused by a compromised control system is determined by the type of process that it controls and by the nature of engineered safety systems (e.g., physical safety mechanisms may be in place that prevents significant damage despite a successful attack on the electronic control systems).

2.6 Detection

The sixth control systems cyber security dimension is Detection. An ideal control systems environment includes detection mechanisms that alert the security group whenever there is an unauthorized event on the control systems. Unauthorized events come in a wide variety of forms and would include activities such as an unauthorized user attempting to gain access to control systems or a counterfeit message from a front end processor to a remote terminal unit (RTU).

2.7 Recovery

The seventh control systems cyber security dimension is Recovery. Ideal control systems can be restored to an uncompromised state immediately after an attack is detected. Recovery time is related to Damage Potential because the cost of a successful attack correlates with the length of time that the control system is in a compromised state. Damage will tend to be less severe if the time to recover is minimized. However, the relationship between Recovery Time and Damage Potential is highly non-linear and highly system dependent.

2.8 Cyber Security Dimensions Drive Cyber Security Metrics

Although perfection may not be feasible for any of these seven dimensions of control systems cyber security, the measurement of how nearly the system approaches the ideal in each dimension is a useful way to think about the cyber security risk to the system. Table 2 is a check list that can assist in consideration of each dimension, and ensure measurement and tracking of system attributes related to each cyber security dimension.

Ten technical security metrics have been defined to support efforts in establishing measures for each control systems cyber security dimension. At least one technical security metric is defined for each control systems cyber security dimension. The ten technical metrics and their application are discussed in the next section.

Table 2. Control Systems Cyber Security Dimension Checklist and Potential Metrics.

Security Dimension	System Attributes to Consider		Recommended metrics
1. Security Group Knowledge	Network diagrams up to date?	☐	Rogue Change Days
	Software change detection?	☐	Security Evaluation Deficiency Count
	Hardware change management?	☐	
	Has the system been scanned for vulnerabilities?	☐	
	Are all communication channels identified?	☐	
2. Attack Group Knowledge	Encrypted traffic?	☐	Data Transmission Exposure
	System specifications removed from control systems?	☐	
	Open source information?	☐	
	Does the system respond to network probes?	☐	
3. Access	Are all allowed network channels needed?	☐	Reachability Count
	Physical barriers to computer access?	☐	Attack Path Depth
	Is authentication required?	☐	
4. Vulnerabilities	Known vulnerabilities?	☐	Known Vulnerability Days
	Are authentication credentials strong?	☐	Password Crack Time
5. Damage Potential	Safety systems in place?	☐	Worst Case Loss
	Are safety systems connected to control systems?	☐	
6. Detection	Antivirus software?	☐	Detection Mechanism Deficiency Count
	IDS?	☐	
7. Recovery	Backups of all critical software?	☐	Restoration Time
	Backups or documentation of all critical hardware configurations?	☐	

3. RECOMMENDED TECHNICAL METRICS

Numerous guides and standards documents (including, for example, references 2, 3, 4, and 5) provide useful guidance for developing cyber security metrics programs. They contain suggested metrics of various types. The recommended technical metrics in this primer are based on the framework described in Section 2 . Each recommended metric was selected by considering measurable system attributes that provide meaningful security representation and relationship to risk for each of the seven cyber security dimensions.

Recommended Technical Metrics:

1. Rogue Change Days
2. Security Evaluation Deficiency Count
3. Data Transmission Exposure
4. Reachability Count
5. Attack Path Depth
6. Known Vulnerability Days
7. Password Crack Time
8. Worst Case Loss
9. Detection Mechanism Deficiency Count
10. Restoration Time

The ten recommended technical metrics are listed in Table 2 and the text box to the right. Each metric is associated with one control systems cyber security dimension and there is at least one metric associated with each of the seven cyber security dimensions. Each defined metric is an answer to the question: "What can be objectively measured on the system that is a reasonable representation of how nearly the system approaches the ideal of its associated control systems cyber security dimension?" As this is a framework, *the metrics chosen may be different, but there should be at least one metric for each of the seven control systems cyber security dimensions*. These ten recommended metrics have been successfully applied to two case studies detailed in Section 4 and are described in sufficient detail to provide a starting point for development of an organization's metrics plan. A control system owner/operator should use these examples to consider how the metrics framework can be applied to their own control systems environment in a consistent manner over time to track progress. Researchers used statistical validation via a comparison with more than two hundred findings from eight control systems cyber security assessments to confirm the coverage and usefulness of the metrics suggested in this primer.[6]

3.1 Rogue Change Days

Rogue Change Days are the number of rogue changes multiplied by the number of days the changes were unknown to the security group. A rogue change is any change to the system configuration without prior notification to the security group. For example, if two modems were added to the control systems environment without the knowledge of the security group and this change was not discovered by the security group until ten days later, this would add 2 * 10 → 20 rogue change days to the metric calculation. This is the first recommended metric for the 'Security Group Knowledge' security dimension. Its' ideal value is zero.

3.2 Security Evaluation Deficiency Count

The Security Evaluation Deficiency Count is the number of control systems network devices that have not undergone a cyber security evaluation. This metric emphasizes the need to measure and track system knowledge about the security attributes of the control systems. For example, if two RTUs that have not undergone security evaluations and one programmable logic controller (PLC) that has undergone security evaluation have been added to the control systems, this would add a count of 3 – 1 → 2 to this

metric calculation. This is the second recommended metric for the 'Security Group Knowledge' security dimension. Its' ideal value is zero.

3.3 Data Transmission Exposure

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 are clearly more sensitive than others; however, to make the metric easier to obtain, it is simply a count of the number of clear text channels used by the control systems. For example, if Telnet is used to connect to the control systems environment from the Internet and if it is the only channel used for external access, then the value of the metric is one. Telnet channels are included in this metric because Telnet uses a clear text protocol that attackers can tap into to obtain passwords and other sensitive data. This is the recommended metric for the 'Attack Group Knowledge' security dimension. Its' ideal value is zero.

3.4 Reachability Count

The Reachability Count is the number of access points in relation 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 is a count of the incoming and outgoing network communication channels plus the number of physical access data channels. For example, the reachability count (from the Internet) for a control system protected by a firewall is calculated as follows. Suppose the control systems environment consists of ten machines with two open TCP/IP ports each, and suppose the firewall prevents access to one of the two ports on each machine, but has no outgoing restrictions. The metric value is ten incoming channels (one for each machine) plus ten outgoing channels (one for each machine), 10 + 10 → 20. This is the first recommended metric for the 'Access' security dimension. Its' ideal value is zero.

3.5 Attack Path Depth

The Attack Path Depth is the minimum number of independent, single-machine compromises required for a successful attack from an external network. This metric emphasizes the importance of having multiple layers of defense. A system configuration that can be successfully attacked by a single exploit should be avoided. For example, the Attack Path Depth metric has a value of one if there is a modem that provides remote access from the public telephone network to critical control systems components, because for this case, a successful attack requires only the compromise of a single device. This is the second recommended metric for the 'Access' security dimension. Its' ideal value is infinity.

3.6 Known Vulnerability Days

The Known Vulnerability Days metric is the sum of known and unpatched vulnerabilities, each multiplied by their exposure time interval. A key assertion is that the longer a vulnerability is known, the greater the risk that it will be exploited. The value of the metric increases each day when there are known and unpatched vulnerabilities. For example, if there are exactly three known and unpatched vulnerabilities on a given system, and if those vulnerabilities were publicly announced two weeks ago today, the current value of the metric should be calculated as 3 * 14 → 42 known vulnerability days. This is the first recommended metric for the 'Vulnerabilities' security dimension. Its' ideal value is zero.

3.7 Password Crack Time

The Password Crack Time is the shortest time (in days) needed to crack a single password for any account on the system. This metric is a measure of the minimum amount of time an attacker would need to compromise the system by password cracking. For example, suppose the encrypted password files have been copied from all of the computers in the control room, and the first of these passwords was cracked in 18 days while the second password was cracked in 30 days using John the Ripper from Openwall.com.[7] If no other passwords were cracked in fewer days, the metric calculation would yield a value of minimum (18, 30) → 18 days. This is the second recommended metric for the 'Vulnerabilities' security dimension. Its' ideal value is infinity.

3.8 Worst Case Loss

Worst Case Loss is the maximum dollar value of the damage/loss that could be inflicted by malicious personnel via a compromised control systems environment. A key assertion is that system risk is strongly related to worst case loss. Although there can be successful attacks where the actual loss is much less than the worst case, a reduction in the worst case loss reduces the potential for loss and, therefore, reduces risk. For an example calculation of this metric, consider a chemical plant in which a major explosion can be triggered by signals from a control system. The value of the metric is the estimated cost resulting from such an explosion in dollars. The estimated cost may include repairs, replacements, and lost revenues from plant downtime. This is the recommended metric for the 'Damage Potential' security dimension. Its' ideal value is zero.

3.9 Detection Mechanism Deficiency Count

The Detection Mechanism Deficiency Count is the number of externally accessible devices that do not have malware detection or 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 potential attacks. For an example calculation of this metric, suppose the control room has fifteen computers each with one or more currently enabled USB ports, and assume that twelve of the computers have antivirus protection installed, but three do not. This would add 15 – 12 → 3 to this metric calculation. This is the recommended metric for the 'Detection' security dimension. Its' ideal value is zero.

3.10 Restoration Time

The Restoration Time is the worst case elapsed time to restore the system to a known uncorrupted state. The metric can be determined by running a test to measure the actual time elapsed from a worst-case compromise to a fully restored and 100% operational system. If a test is not feasible and there have been no cyber security events on the control systems where the restoration time was tracked, the metric may be estimated. For example, assume a situation where all twenty computers in the control room have been compromised by a virus. But, the effects of the virus are relatively benign, allowing the response team to address one computer at a time. In this scenario, individual computers are taken off the network, while the rest of the system operates in a degraded mode. The team cleans the virus from each machine, and then reintroduces the computer to the network and restores the applications in an up-to-date status. If this activity for a single machine takes 1-1/2 hours, the restoration time would yield a metric value of 20 * 90 → 1,800 minutes. This is the recommended technical metric for the 'Recovery' security dimension. The ideal value for the restoration time technical metric is zero.

3.11 Specifications of Recommended Technical Metrics

Detailed metric specifications for each of the ten recommended core technical metrics are listed in Tables 3 through 12.

Table 3. Rogue Change Days metric specification.

Metric Name	Rogue Change Days
Cyber Security Dimension	Security Group Knowledge
Brief Description	The number of rogue changes multiplied by the number of days the changes were unknown to the security group. A Rogue Change is any change to a control system's configuration without prior notification to the security group.
Associated Principle	The system configuration should not be changed without the security group's knowledge.
Key Assertion	The cyber security risk from changes to the system without notification to the security group is, on average, worse than for changes that are announced in a well managed system.
Units	RCD (Rogue Change Days)
Range of Metric	Integer, 0 to infinity (best to worst)
Strengths	A valid measure of the quantity of uncontrolled changes.
Weaknesses	The metric does not include any measure of the actual cyber security impact of changes.
What Data is Collected	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, Human Machine Interfaces (HMI), network devices attached to the local network, and critical computer files, such as programs and graphical display pages.
How Data is Collected	One method for base lining computer files is to simply save a list of file names with associated file dates. A rogue change is then detected when there is a difference between the actual file date and the saved file date. Since file dates might be faked by an attacker, a more secure method is to apply a hash function to the content of each computer file to calculate an identifier that can be saved with each file name and that can be used to determine when the file has been changed.
Tools	Tripwire[8] is an example of a tool that supports change auditing.
Data Collection Frequency	Because the data collection time interval is included in the calculation of the metric, there is built-in incentive to measure often. It is recommended that the data be collected at least monthly.
Discussion	RCD is somewhat conservative because the metric assumes the rogue change was introduced immediately after the previous measurement. Awareness of all devices, components, processes, and topology of a system is an important goal for effective security management. This metric is practical for most systems only after the deployment of tools that support strict change management. Therefore, an automated mechanism that allows management to verify whether the current real system configuration matches the baseline is recommended. Tools of this type are available and feasible even though they have historically not been widely used in control systems to date.

Table 4. Security Evaluation Deficiency Count metric specification.

Metric Name	**Security Evaluation Deficiency Count**
Cyber Security Dimension	Security Group Knowledge
Brief Description	The number of control systems network devices that have not undergone a security evaluation.
Associated Principle	A control system and its components should be evaluated and monitored for vulnerabilities.
Key Assertion	Security evaluation of the network components identifies vulnerabilities and leads to improved cyber security of control systems
Units	Count.
Range of Metric	Integer, 0 to infinity (best to worst).
Strengths	It is important to have some measure of how well control systems security has been evaluated. This metric provides a simple measure.
Weaknesses	The measurement requires technical judgment (not completely objective). The metric does not address variances in risk and potential impact associated with individual devices.
What Data is Collected	Identify all devices connected to the control systems network. Identify which of those devices have been evaluated from a cyber security perspective.
How Data is Collected	First, identify the system boundary. Then, identify all devices (machines) connected to the network within that boundary. Count the devices within that boundary that have not undergone a security evaluation.
Tools	Control systems' configuration diagrams. Piping and instrument diagrams.
Data Collection Frequency	At least annually.
Discussion	Unresolved questions: Do security evaluations become obsolete with the passage of time? If so, how much time makes them obsolete? When a network device is changed significantly, the previous security evaluation should be obsolete. How does an organization objectively determine whether a given device has undergone a sufficient security evaluation? Even though this metric has unresolved questions and requires some subjective judgment, it is important enough to remain in the list of core metrics.

Table 5. Data Transmission Exposure metric specification.

Metric Name	Data Transmission Exposure
Cyber Security Dimension	Attack Group Knowledge
Brief Description	Unencrypted data transmission volume. The number of unencrypted communication channels in use by devices within the control system boundary.
Associated Principle	The system should send no unencrypted information through networks that may be monitored by potential attackers.
Key Assertion	As the quantity of clear text data transmission is reduced, the risk of disclosing sensitive data to potential attackers is also reduced.
Units	Count.
Range of Metric	Integer, 0 to infinity (best to worst).
Strengths	Can be objectively measured.
Weaknesses	Some data are clearly more sensitive than others, but, in order to be objective, this metric only measures the potential for the interception of sensitive data rather than attempting to measure the volume of sensitive data. The metric does not include any measure of data previously made public that may be stored at an external location, such as a search engine cache. Encrypted communication channels may limit the functionality of network monitoring controls.
What Data is Collected	Select one or more appropriate network access points for monitoring. At least one access point should be outside the control systems network perimeter. Points that may be accessible to potential attackers are of the most interest. At each chosen access point, monitor data transmissions during a long enough time to capture all the various types of traffic that occur during normal operations. Count the number of unencrypted channels that terminate on a control systems component. For a TCP/IP network, count the number of IP-address/TCP-port pairs found in the network traffic such that the TCP-port is associated with an unencrypted service and the IP-address belongs to a control systems component. Any observed encrypted communication channels should be authorized by the organization. Unauthorized encrypted communications may represent malicious activity. File Transfer Protocol (FTP) and telnet are examples of unencrypted services that would be included in this count.
How Data is Collected	A network monitor is attached to each network access point selected within the control systems environment.
Tools	Passive network monitoring tools.
Data Collection Frequency	Collect data whenever there is a configuration change that may affect network traffic. It should be collected regularly (monthly, if possible), even if there are no known changes, to verify that only authorized channels are still in use.
Discussion	Any data transmitted to external networks is potentially useful to an attacker. Ideally, there should be no response to external network probes and no “clear text” transmissions. This is a conservative metric because there is no attempt to determine the sensitivity of the data. All data has the potential to provide useful information to an attacker; therefore, all data is treated as equally sensitive. Some network paths are more critical than others, but during a multistage attack, an attacker may gain access to an internal network by first penetrating the system through an external network path. This metric does not include any measure of the amount or nature of data that is made available to potential attackers in response to probes, but vulnerability scanners often identify such behaviors.

Table 6. Reachability Count metric specification.

Metric Name	Reachability Count
Cyber Security Dimension	Access
Brief Description	Number of access points (relative to a specific point of origin, such as the Internet).
Associated Principle	Number of communication channels associated with control systems components should be minimized, including network connections, TCP/IP ports/services, physical access to USB ports, and portable storage media drives.
Key Assertion	Reduction in the number of access points reduces cyber security risk.
Units	Count.
Range of Metric	Integer, 0 to infinity (best to worst).
Strengths	Objectively measureable by network scanners and router/firewall rules analysis. Any communication channel has the potential to become part of an attack path. This metric helps to identify potential attack paths.
Weaknesses	The relationship between this metric and level of cyber security is unknown. Any non-zero value implies a security risk, but the risk does not increase linearly as the metric increases.
What Data is Collected	Network configuration and physical access information.
How Data is Collected	If network scanning does not affect operations, use network scanners and network configuration information to identify all network communication paths, including open TCP/IP ports. Some control systems environments are adversely affected by scanning; therefore, scanning should only be done when the operational impact is acceptable. To avoid adverse affects on operations, scan representative machines off-line and obtain the metric by extrapolation and analysis of router/firewall rules. The identification of physical access ports must be determined by inspection.
Tools	Network configuration discovery tools such as Nmap. Passive network traffic data collection tools can also be useful, but passive tools do not identify communication paths unless they happen to be transmitting during the data collection period. For example, an open TCP/IP port with no active clients would not be identified with a passive tool.
Data Collection Frequency	Data collection frequency depends on the stability of the system configuration and should occur whenever there has been a known configuration change that affects accessibility.

Table 6. (continued).

Metric Name	Reachability Count
Discussion	This metric applies to a specific point of origin. The Reachability Count is the sum of three components: (1) the number of ports/services that can be reached from the point of origin, (2) the number of active physical network ports with outgoing network connectivity to the point of origin, and (3) the number of machines with unrestricted portable storage media drives, including unrestricted USB ports accessible from the point of origin. "Outgoing network connectivity" means the network configuration allows the port to originate two-way connection-oriented sessions to some machine located at the point of origin. The system manager may choose to combine the network and physical reachability counts or track them separately. The external network of most concern is the Internet. However, all external networks should be treated as non-trusted. Because of the possibility of penetration of the perimeter, the system manager may choose to also calculate this metric for 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 perimeter or some other partially controlled area or combination of areas within the fence as defined by the system 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, physically locked, or blocked with glue • Host-based or device-based port encryption • Ports restricted by end-point control software.

Table 7. Attack Path Depth metric specification.

Metric Name	**Attack Path Depth**
Cyber Security Dimension	Access
Brief Description	Minimum number of independent, single-machine compromises required for a successful attack from an external network.
Associated Principles	Compartmentalization. Defense in depth.
Key Assertion	Risk is reduced by increasing the number of independent steps required for a successful attack.
Units	Count.
Range of Metric	Integer, 0 to infinity (worst to best).
Strengths	This metric emphasizes the importance of having multiple layers of defense.
Weaknesses	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.
What Data is Collected	Identify all the network connections that form an unbroken set of links between the control system's critical components and the external networks. Identify the path that traverses the fewest number of computers.
How Data is Collected	If network scanning is allowed, use network scanners and network configuration information to identify all network communication paths, including open TCP/UDP ports. If scanning of the entire network is not allowed, scan representative machines off-line and obtain the metric by extrapolation and analysis of router/firewall rules.
Tools	Network configuration discovery tools such as Nmap.
Data Collection Frequency	Data collection frequency depends on the stability of the system configuration and should occur whenever there has been a known configuration change that affects accessibility.
Discussion	This metric encourages network configurations that have several layers of protection.

Table 8. Known Vulnerability Days metric specification.

Metric Name	Known Vulnerability Days
Cyber Security Dimension	Vulnerabilities
Brief Description	The sum of known and unpatched vulnerabilities, each multiplied by exposure time interval. (Note: Along with other vulnerabilities, this measure also includes all known vulnerabilities in firmware which haven't been fixed.)
Associated Principle	The time between vulnerability discovery and repair should be small.
Key Assertion	The longer a vulnerability is known, the greater the risk that it will be exploited.
Units	Vulnerability Days.
Range of Metric	Integer, 0 to infinity (best to worst).
Strengths	The metric is objective, measurable, and encourages the patching of known vulnerabilities.
Weaknesses	The metric does not include any measure of the vulnerability potential or of "zero-day" vulnerabilities and exploits that are known to attackers but are not publically disclosed. The metric depends on the strength of the vulnerability identification tools being used. Also, the metric does not address the severity rating of the vulnerabilities.
What Data is Collected	Known vulnerabilities that apply to components of the control system and the discovery or public disclosure date for each known vulnerability.
How Data is Collected	Vulnerabilities that apply to the system are identified by vulnerability identification tools and by comparing system components to the components associated with publicly disclosed vulnerabilities. For publicly disclosed vulnerabilities, the discovery time is the disclosure date from the public vulnerability database. For vulnerabilities that are discovered locally, such as configuration errors, the local discovery date is applicable.
Tools	Nessus and other vulnerability discovery tools.[9,10] Note: Control systems environments may be adversely affected by vulnerability discovery tools; therefore, these tools should not be used to target critical components of operational systems. The tools may be used to target representative machines off-line or non-critical components when the potential impact on operations is minimal. Public vulnerability databases (e.g., Common Vulnerabilities and Exposures [CVE]). Vulnerability prioritization tools, such as CVSS.[11]
Data Collection Frequency	Data should be collected whenever there is a change in known vulnerability status. The system should also be scanned for vulnerabilities often, but within operational constraints that avoid the danger of impacting critical operational components (suggested weekly or when there is a known configuration change). Public vulnerability databases should be checked regularly and often (suggested daily). The discovery of a vulnerability that applies to the system or the repair/mitigation of a vulnerability are events that change vulnerability status, and therefore affect the value of the metric. The calculation of the metric should be done at regular intervals (suggested daily or weekly) because the metric changes with the passage of time even if there is no change in the status of the known vulnerabilities.
Discussion	The exposure time interval for each known and unpatched vulnerability is the current date minus the date on which the vulnerability was discovered and/or made public. This metric is affected by vulnerability discovery rate and patch rate. Vulnerabilities may result from design errors, implementation errors, and from misconfigurations, such as inappropriate trusted relationships between machines. Some vulnerabilities are more significant than others. Tools such as CVSS can be used to determine priority categories for all known vulnerabilities and then the metric can be applied to each category.

Table 9. Password Crack Time metric specification.

Metric Name	Password Crack Time
Cyber Security Dimension	Vulnerabilities
Brief Description	This metric is the minimum time (in days) needed to crack a single password for any account on the control system.
Associated Principle	Credential keys should be strong and should be changed regularly.
Key Assertion	Control systems cyber security improves as password crack time increases.
Units	Days.
Range of Metric	Real number, 0 to infinity (worst to best).
Strengths	A valid measure of the time an attacker would need to compromise the system by password cracking. Passwords are the most common form of authentication.
Weaknesses	Does not measure the strength of other authentication mechanisms. Does not account for how well the hashed password data is protected.
What Data is Collected	Data collected for this metric is the encrypted password files (hashes) from all machines on the system. For example, all password files from UNIX servers, configuration data for web servers, database servers, Windows® workstations, and control systems Human Machine Interface (HMI) machines. A password cracking tool is then applied to each password file instance.
How Data is Collected	Copy account data for each instance of an operating system and each instance of an application that implements authentication. For UNIX type machines, the password hash file or shadow file is simply copied. For Windows® machines, a password hash file is created from registry data by using the Pwdump[12] freeware. The password cracking operation should be conducted on a system configured with strong security controls.
Tools	Password cracking tools are available commercially and for free download. Examples include L0phtCrack[13] and John the Ripper.[7] Pwdump is freeware used to obtain password hash data from Windows® machines.
Data Collection Frequency	Data should be collected whenever passwords change. Password cracking software should run continuously until there are new hash data or until all passwords are cracked.
Discussion	This metric is an important measure because passwords are by far the most common form of authentication. The older a password becomes, the more likely it is to be compromised; therefore, the system manager may choose to track password age as well as password crack 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 can crack very weak passwords virtually instantly. Passwords used for authentication at the perimeter are particularly important and, therefore, perhaps crack time for those passwords should be measured separately from other passwords used on the system. The security group should ensure that vendor supplied passwords and passwords commonly used by maintenance personnel are included in the password cracker's dictionary.

Table 10. Worst Case Loss metric specification.

Metric Name	Worst Case Loss
Cyber Security Dimension	Damage Potential
Brief Description	The maximum dollar value of the damage/loss that could be inflicted by malicious personnel via a compromised control system.
Associated Principle	Mechanisms that are independent of the control systems should provide protection such that the cost of damage due to control systems malicious behavior is minimized.
Key Assertion	Risk is strongly related to the potential for loss.
Units	Dollars.
Range of Metric	Integer, 0 to infinity (best to worst).
Strengths	This metric strongly correlates to risk.
Weaknesses	An objective measurement is usually not feasible. Estimation by careful analysis is required.
What Data is Collected	Control systems interfaces to the process under control. Detailed process and plant design and configuration information.
How Data is Collected	Facility documentation and expert elicitation.
Tools	Standard safety analysis methods apply.
Data Collection Frequency	The metric should be reassessed whenever there is a system configuration change that may impact the potential worst case loss.
Discussion	Although there can be successful attacks where the actual loss is much less than the worst case, 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 this question: "If the control system is under the control of a malicious person, what damage can be done?" If safety systems are not completely independent of the control systems (i.e., a safety system connected to the control systems network), then safety systems may also be compromised by an attacker; therefore, it should not necessarily be assumed that such a safety system will prevent damage when estimating the worst case loss.

Table 11. Detection Mechanism Deficiency Count metric specification.

Metric Name	**Detection Mechanism Deficiency Count**
Cyber Security Dimension	Detection
Brief Description	The number of externally accessible devices without any malware/attack detection mechanism.
Associated Principle	The system should be constantly monitored for malicious behavior and alarms should be raised when detected.
Key Assertion	Detection mechanisms reduce risk, especially when applied to devices that are potential entry points for attacks.
Units	Count.
Range of Metric	Integer, 0 to infinity (best to worst).
Strengths	Easy to obtain.
Weaknesses	The effectiveness of detection mechanisms is not addressed by this metric. Correlation to risk is unknown.
What Data is Collected	Identify which devices are accessible externally. Identify the location of each detection mechanism (Antivirus Software, IDS, etc.).
How Data is Collected	System configuration information.
Tools	Not applicable.
Data Collection Frequency	Whenever the configuration changes.
Discussion	Although detection mechanisms such as Antivirus Software and IDS cannot be expected to detect every attack, they do detect some attacks and, therefore, should be employed at network entry points. This metric encourages the use of these mechanisms.

Table 12. Restoration Time metric specification.

Metric Name	Restoration Time
Cyber Security Dimension	Recovery
Brief Description	Worst case elapsed time to restore the control system to normal business operations in a known uncorrupted state.
Associated Principle	The time needed to restore the system to an uncorrupted state should be minimal.
Key Assertion	Risk is correlated with the amount of time needed to restore the control system to an uncompromised state.
Units	Minutes.
Range of Metric	Real number, 0 to infinity (best to worst).
Strengths	Simple model.
Weaknesses	On some systems it may be difficult to measure accurately. Many systems require continual operation and there may be no known incidents that included full system restoration.
What Data is Collected	Measure the actual time elapsed from "start" to a fully restored and 100% operational system.
How Data is Collected	Record elapsed time during the recovery process. This data should be collected for actual cyber security events if they have occurred. If there have been no recorded cyber security incidents, a test of the recovery process would include restoration of all software.
Tools	Control systems themselves provide information characterizing interruptions in normal operations and restoration of normal operation. System and external clocks provide the measurement of time.
Data Collection Frequency	As needed.
Discussion	The metric can be determined by running a test to measure the actual time elapsed from the "start" of an interruption of operations to a fully restored and 100% operational system. If it is impractical to perform that kind of test on an operational system, then this data should be collected for actual cyber security events should they occur. If a recovery test is not feasible, then a worst case recovery analysis may be used to estimate recovery time. Each organization must establish a restoration time objective for critical business processes. The target value for restoration time depends on the associated processes and environments. For processes with high-availability requirements, a target value near zero may be appropriate. Some environments may have less demanding recovery requirements.

4. CASE STUDIES

The recommended technical metrics were applied in two case studies of operating control systems. These studies help to identify the practical limitations associated with data collection and provide specific examples of how the metrics could be obtained and used.

Case Studies of Operating Control Systems:

1. Distributed control system for a chemical processing plant
2. SCADA System used for power distribution

4.1 Case Study 1

The first case study was 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 DCS 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 DCS with an access control list so that only the incoming TCP/IP connections with origination addresses that match the access control list are allowed through the router. The system boundary is defined to be the boundary that encloses the processing plant and the DCS networks that are within the control room. A dedicated telephone line connects the control system to the plant data archive system, which has direct connectivity to the corporate network. The corporate network affects the security of the DCS, but, for this study, the corporate network was not considered to be part of the system.

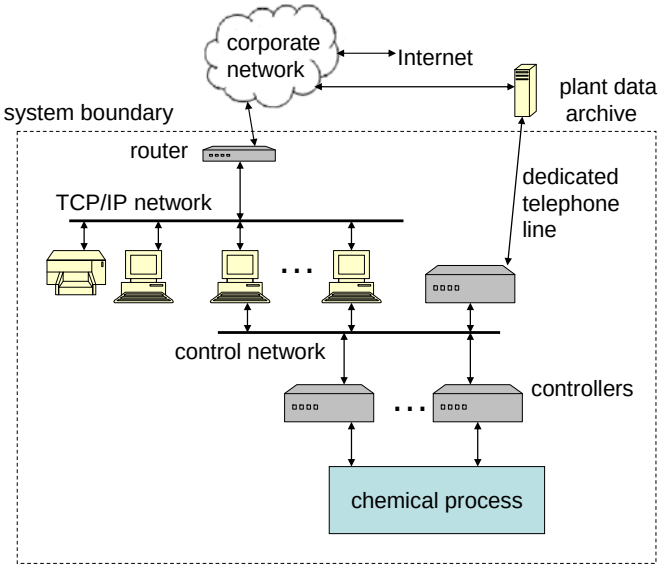

Figure 1. Simplified network diagram of Case Study 1 (DCS).

The DCS for this case study consists of a TCP/IP network that provides connectivity for eleven workstations and two printers, and a proprietary control system network that provides connectivity to approximately thirty 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.

The metric values obtained for Study 1 are shown in Table 13, which also show the ideal value and suggested target value for each metric. The suggested target value is the value that could be obtained by changing the system configuration to improve cyber security while retaining required functionality. This case study verified that the values of the metrics listed in Table 13 could be obtained using tools that are readily available.

Table 13. Case Study 1 metric values.

Metric Name	Metric Value	Ideal Metric Value	Suggested Target Value
Rogue Change Days	0	0	0
Data Transmission Exposure	23	0	1
Reachability Count (N_T)[a]	164	0	1
• Physical (N_p)	2	0	0
• Services (N_S)	149	0	1
• Outgoing (N_o)	13	0	0
Attack Path Depth	2	∞	4
Known Vulnerability Days (high priority)	16,416 vuln. days	0	0
Known Vulnerability Days (low priority)	15,877 vuln. days	0	0
Password Crack Time	> 30 days	∞	>30 days
Worst Case Loss	$100 M	$0	unknown
Detection Mechanism Deficiency Count	12	0	0
Restoration Time	120 minutes	0	120 minutes

a. $N_T = N_p + N_S + N_o$

4.2 Case Study 2

Case Study 2 was for a power distribution Supervisory Control and Data Acquisition (SCADA) system. Figure 2 is a simplified network diagram of the SCADA system showing the main components and connectivity. The SCADA controls seven power distribution substations. There is one PLC and one Voice over Internet Protocol (VoIP) phone physically located at each of the seven substations and twenty-five electric power meters connected directly to the network. One engineering workstation (EWS), two front-end processors (FEP), three HMI hosts, one printer, and the SCADA firewall are all physically located in the main substation control room. All devices are logically connected to a single TCP/IP network through routers and switches. The only connection from the SCADA network to the external network is through the SCADA firewall with a Virtual Private Network (VPN).

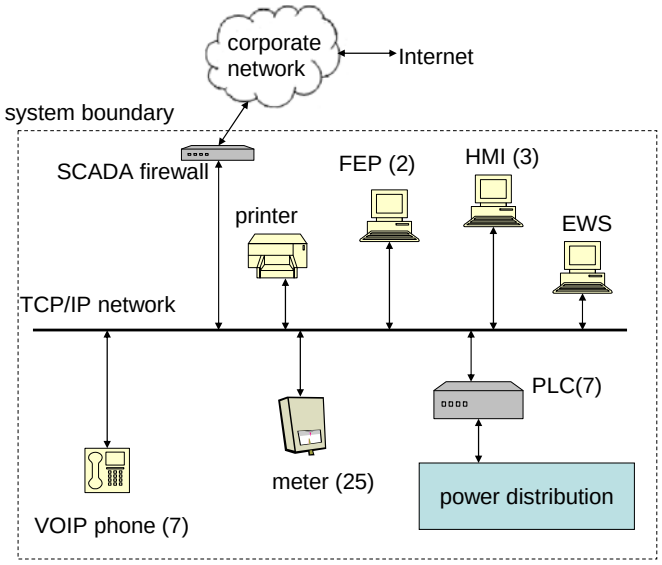

Figure 2. Simplified network diagram of Case Study 2 (SCADA).

Data were collected by testing on a duplicate laboratory system when feasible, rather than the live SCADA system, to avoid potential disruption of operations. For the cases where the laboratory system did not provide sufficient information, data was also collected on the live operating SCADA system.

All of the passwords for this case study were cracked in less than one day using John the Ripper. The password hashes were "LAN manager" hashes that are known to be easy to crack. The metric values obtained for Case Study 2 are shown in Table 14.

Table 14 also gives the ideal value for each metric and the suggested target value. The suggested target value is the estimated value that could be obtained by changing the system configuration as suggested by the comments in the table.

Table 14. Case Study 2 metric values.

Metric Name	Metric Value	Ideal Metric Value	Suggested Target Value	Comment
Rogue Change Days	unknown	0	0	Mechanisms are needed to detect rogue changes
Data Transmission Exposure (Internet)	0	0	0	No unencrypted protocols are used through the corporate firewall
Data Transmission Exposure (Intranet)	28	0	0	Avoid using unencrypted protocols through the control system firewall.
1. Reachability Count (external) (N_T)[a]	0	0	0	External reachability count total.
1. Physical (N_p)	0	0	0	No physical access outside the control room.
1. Services (N_S)	0	0	0	External connectivity only through a VPN.
1. Outgoing (N_o)	0	0	0	Firewall prevents outgoing connections.
2. Reachability Count (local) (N_T)[a]	91	0	1	Local Reachability count total from within control room and/or from VPN.
2. Physical (N_p)	4	0	0	HMI and EWS have unrestricted physical access USB ports inside the control room.
2. Services (N_S)	87	0	1	Suggest further firewall rules to restrict VPN access.
2. Outgoing (N_o)	0	0	0	Firewall prevents outgoing connections.
Attack Path Depth	1	∞	2	Suggest further firewall rules to restrict VPN access to increase attack path depth.
Known Vulnerability Days (high priority)[†]	36,318 vuln. days	0	0	Recommend the operating system be patched to current level. The PLC vulnerability has no available fix but PLC reachability improvement could mitigate the vulnerability.
Known Vulnerability Days (low priority)[b]	18,624 vuln. days	0	0	All known vulnerabilities should be mitigated if feasible.
Password Crack Time	0 days	∞	30+ days	Use longer passwords and avoid LAN manager hash of the password from being stored on Windows® machines (see http://support.microsoft.com/kb/299656).
Worst Case Loss	unknown	$0	unknown	Worst case power outage from a cyber attack is estimated to be 6 hours duration. The dollar value of such an outage has not been estimated.
Detection Mechanism Deficiency Count	46	0	0	There are no detection mechanisms on the system. The addition of AV software and firewall restrictions can reduce metric value to zero.
Restoration Time	72 hours	0	5 hours	Worst case restoration time is high because it requires a complete software rebuild.

a. NT = Np + NS + No.

b. High priority vulnerabilities allow penetration and/or CVSS score is greater than 7.

5. SUMMARY

The control systems cyber security framework consists of seven control systems cyber security dimensions, each related to risk. Case studies and reviews of control systems cyber security assessments have demonstrated the framework's ability to address control systems risk exposure. As a result, the seven control systems cyber security dimensions are recommended as a foundation for managing cyber security of control systems environments and provide a framework for the ten recommended technical metrics.

The ten recommended technical metrics support assessment of cyber security risk exposure over time. The recommended metrics have been applied to control systems environments and proven to be practical and useful. However, every system and facility is unique, so there may be a need to select tailored metrics or measurement technologies in line with particular circumstances. An organization's tailored technical metrics should have at least one metric for each of the seven cyber security dimensions.

An important use of the recommended technical metrics is in tracking the improvement or degradation of control systems cyber security posture along all seven dimensions of cyber security. As the cyber security posture improves, the risk to control systems from a cyber attack diminishes. Diligent use of the control systems cyber security framework and application of the technical metrics will aid in making more effective cyber security decisions for control systems environments.

6. REFERENCES

1. A. Jacquith, *Security Metrics*, Addison Wesley, 2007.
2. Chemical Sector Cyber Security Program (CSCSP), *Guidance for Addressing Cyber Security in the Chemical Industry*, Technical Report, CSCSP, May 2006.
3. E. Chew, A. Clay, J. Hash, N. Bartol, and A. Brown, Guide for Developing Performance Metrics for Information Security, NIST Special Publication 800-80, May 2006.
4. R. Ross, S. Katzke, A. Johnson, M. Swanson, and 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.
5. M. Swanson, N. Bartol, J. Sabato, J. Hash, and L. Graffo, *Security Metrics Guide for Information Technology Systems*, NIST Special Publication 800-55, National Institute of Standards and Technology (NIST), Gaithersburg, Maryland, July 2003.
6. M. McQueen, W. Boyer, S. McBride, M. Farrar, and Z. Tudor, "Measurable Control System Security through Ideal Driven Technical Metrics", S4: SCADA Security Scientific Symposium, January 23, 2008.
7. John the Ripper (http://www.openwall.com).
8. Tripwire (http://www.tripwire.com/products/index.cfm).
9. Nessus vulnerability scanner (http://www.nessus.org).
10. Metasploit (http://www.metasploit.com).
11. M. Schiffman, *A Complete Guide to the Common Vulnerability Scoring System (CVSS)*, https://www.first.org/cvss/cvss-guide.html, Technical Report, Forum for Incident Response and Security Teams (FIRST), published June 7, 2005, Web site visited July 11, 2006.
12. Pwdump (http://www.darknet.org.uk/2006/10/download-pwdump-142-and-fgdump-134-windows-password-dumping).
13. LC5 (http://www.securityfocus.com/tools/1005).