

Cyber Security Assessments of Industrial Control Systems

November 2010



**Homeland
Security**

CPNI

Centre for the Protection
of National Infrastructure

Control Systems Security Program

National Cyber Security Division



CPNI

Centre for the Protection
Of National Infrastructure

Executive summary

Cyber security has become a vital part of conducting business in today's world. The threats to organisations and individuals are real. Industrial Control Systems (ICSs) were originally built as stand-alone systems that were not interconnected and had little in the way of security protections. The internet and ubiquitous internet protocol networks have changed the design of many ICS such that the control network is now often a protected extension of the corporate network. This means that these delicate ICSs are potentially reachable from the Internet by malicious and skilled adversaries.

One tool that an ICS asset owner may utilise to assess the risk to the ICS is to procure and facilitate a cyber security assessment. The ICS cyber security assessment identifies and seeks to mitigate vulnerabilities that would allow an attacker to disrupt or take control of the system. Many considerations have to be taken into account because of significant differences between an ICS cyber security assessment and the tests that would be performed in a standard corporate environment. For example, several tools employed in such a test could have a serious impact on the ICS itself. Various ICSs will malfunction or halt completely when security tools, such as scanners, are run on the network. Therefore, the asset owner and assessment team must understand the potential implications of testing on a production system. Whenever possible, cyber security tests should be performed on a backup or offline ICS.

This guide aims to assist asset owners to maximise the return on their investment when commissioning assessments of their ICSs.

The guide provides an overview of the assessment process so users understand how to execute an ICS cyber security assessment. This guide also covers the process of planning an ICS cyber security assessment, including how to select testing areas. The test plan specifies the correct amount of detail to meet the needs of the asset owner while retaining the flexibility to use all the skills of the assessment team. The details of the actual testing process in this guide familiarise the asset owner with the steps and reasons behind the testing process. The reporting process for an ICS cyber security assessment is also covered in this guide.

In addition to explaining actual security testing, the pros and cons of a number of alternate vulnerability testing methods for ICSs are also considered so tests can be tailored to the specifics of the ICS and needs of the organisation.

The best assessment methodology is the one that promises the highest vulnerability reduction at lowest cost. The benefit from a vulnerability assessment is proportional to the number of vulnerabilities that are identified for remediation. The actual benefit is the decreased risk due to vulnerability remediation. The benefit is therefore dependent on the asset owner's ability to mitigate the identified vulnerabilities. The asset owner should ensure that the assessment team provides adequate vulnerability details and mitigation information for the ICS administrators or vendors to efficiently and effectively remediate each security weakness. Collaboration between assessment and ICS personnel throughout the assessment allows knowledge transfer both directions and efficient assessment and mitigation performance. ICS staff can gain security knowledge directly applicable to their system and mitigate vulnerabilities as they are identified.

Contents

Disclaimers	4
ICS Assessment versus a typical IT penetration test.....	5
Types of cyber security testing.....	5
Levels of disclosure	7
Focus of testing	8
Impacts of testing.....	9
Testing process overview.....	11
Overview.....	11
Choosing the assessment team.....	13
The test plan.....	15
Selecting the attack vectors	15
Assessment execution.....	19
Assessment reporting	25
Report executive summary	25
Report introduction	25
Report target(s)	27
Attack scenarios	33
Network assessment	34
Report conclusion.....	34
Vulnerability mitigation and vendor engagement.....	34
Assessment variables	36
Alternative methodologies.....	38
Laboratory assessment.....	38
Production system	41
End-to-end penetration assessment	42
Component testing.....	44
Technical documentation review	47
Functionality and configuration review	49
Staff interviews	51
Risk assessment.....	52
Assessment methodologies conclusion	55
Glossary.....	63
Acronyms.....	63
Nomenclature	65
Figures and Tables	67

Overview

This guide has been prepared to assist asset owners in procuring and executing cyber security tests of their Industrial Control, Supervisory Control and Data Acquisition (SCADA), Distributed Control (DCS) and/or process control (PCS) systems, hereafter generically referred to as an industrial control system (ICS). The guide's purpose is to educate asset owners on the general process of a cyber security test and provide insight on specific testing methods so owners learn to prescribe a custom assessment that will maximise the output of their testing budget.

This guide also doubles as a checklist for internal teams performing cyber security assessments to ensure their plans cover the high-risk areas of an ICS. It lists some possible testing methods and describes pros and cons for each method based on the cyber security ICS testing experience of Idaho National Laboratory (INL). Asset owners are able to apply this information in the decision-making process for planning an ICS assessment.

This guide does not describe how to execute specific cyber security tests; rather, it focuses on what should be covered in an ICS cyber security assessment. General cyber security guidance can be followed to meet the operational and security goals of individual ICS components, but like other computer networks, the security goals, threats and potential impacts vary between systems. For this reason, cyber security guidance cannot become prescriptive. Security standards and best practices must be used as guidelines, tailored to the individual system's requirements. Although the ICS domain has many traits in common with the corporate IT domain, the security goals and potential consequences of an attack are very different. This document focuses on the security goals and risks common to the ICS domain and how it interacts with the rest of the network.

The authors prepared this guide under the assumption that the reader has a general understanding of ICSs. For this reason, the guide does not cover best practice topics on the way that ICSs are designed and used.

Disclaimers

A secure ICS does not exist, which means that hidden vulnerabilities are still possible in an ICS, even after a clean report from a cyber security assessment. Cyber security should be perceived as a process rather than a project. A cyber security assessment of an ICS is viewed as a snapshot in time. An ICS needs to be iteratively tested, based on triggers such as changes to the system or an elapsed period of time. One reason for repeated testing is that most ICSs are built using commercial off-the-shelf hardware and software. New vulnerabilities often are discovered in the current operating systems and third-party software that make up today's ICSs. The implications of these vulnerabilities to the ICS domain may not be obvious, but could be exposed by a cyber security assessment. Also, one assessment team may have skills or ideas that uncover problems that another team missed in previous tests. New exploit and mitigation techniques are continually developed, so additional findings and mitigation recommendations should be expected from subsequent vulnerability assessments.

This guide considers several cyber security tools and software programs. These references serve as examples rather than endorsements. For every tool referenced, other proprietary and open source alternatives may exist which implement the same features with varying levels of effectiveness.

Cyber security testing activities may have adverse effects on any target system, but especially on an ICS. Cyber security tests often employ port and vulnerability scanners that make rapid requests to an Internet Protocol (IP) address, often with invalid data. These scans alone often cause a victim process or entire machine to fail. When the target is an active ICS server, this failure could have serious and drastic consequences. All cyber security testing should be well planned and communicated with the equipment owners and operators so that potential faults are resolved or mitigated. The testing methods presented in this document are, therefore, to be employed at the asset owner's own risk.

ICS Assessment versus a typical IT penetration test

Although similarities exist in the tools and methodologies used, an ICS cyber security assessment differs significantly from an IT penetration test. Some of these differences concern the goals, focus and impact of testing.

Types of cyber security testing

The goals of testing can generally be described as assessing the level of security and/or identifying vulnerabilities for remediation/mitigation. Vulnerabilities can be identified by attacking the system as a hacker would or by evaluating the system.

A vulnerability assessment simply identifies and reports noted vulnerabilities and security weaknesses in the target system. The assessment team generally reviews code, settings, etc. for known security weaknesses. Many security tools and techniques used by penetration testers and hackers are used to help identify and validate vulnerabilities. The customer may specify the level of vulnerability verification. For example, practices known to lead to vulnerabilities can be identified for remediation to decrease assessment costs, increase vulnerability identification coverage, and maximise the security of the system.

A penetration test attempts to duplicate the actions of an attacker. The goal of external penetration testing is to find weaknesses in the company's network that could allow an attacker to access the enterprise environment from the Internet. Internal testing attempts to find and exploit vulnerabilities to determine whether unauthorised access or other malicious activity is possible from inside the target network.^a

This can give an indication of the system's ability to withstand attack originating at the location the test team is given access, not including any components that were defined as off limits. Vulnerabilities that were exploited to meet the objective of the test will be identified, but this method does not identify a high percentage of vulnerabilities. This form of security testing is used to answer the question: Can an attacker achieve the identified actions given the access they were granted (potentially no access)?

The company's security team can be tested while gaining experience by actively defending against penetration testers in red team exercises. Red team exercises have the goals of improved readiness of the organisation, better training for defensive practitioners and inspection of current performance levels. Independent red teams can provide valuable objectivity regarding both the existence of vulnerabilities and the efficacy of defences and mitigating controls already in place and even those planned for future implementation.^b As in a real attack, testers attempt to conceal their actions and most corporate personnel are not given advanced notification of the test. Red team exercises may include social engineering attacks, which are attempts to trick employees into divulging security information. Testers may call unsuspecting employees pretending to be someone in a position of authority, and convincing the trusting employee to divulge information later used to infiltrate into the system. Social engineering tests may include phishing attacks, where the tester sends legitimate looking e-mails to employees requesting information or containing links to malicious websites.

^awww.pcisecuritystandards.org/minisite/en/docs/information_supplement_11.3.pdf

^b www.sans.org/critical-security-controls/control.php?id=17

Red team exercises can be valuable practice for ICS administrators because vulnerabilities in ICS products cannot be fully mitigated with perimeter protection. IDS signatures can be tailored to identify invalid or abnormal network traffic, but network administrators must be able to respond quickly and appropriately in order to halt the potential attack without impairing critical ICS functions.

Ideally, an application, component, or network will be secured using vulnerability assessments and then validated by penetration testing. A system should first be iteratively assessed for secure practices by an internal security expert or team with identified vulnerabilities remediated until it has reached an acceptable level. This should be repeated using an external assessment team. Using internal and external assessment teams increases the coverage of identified vulnerability types. Security can then be tested by a penetration team and/or red team exercises.

An ICS cyber security assessment should be a collaborative effort between the assessment team, asset owner and sometimes the vendor. The assessment team is provided detailed drawings of the system in advance, along with network device information such as firewall rules and switch and router configurations. Vendors may make the source code for their applications available to the team. No attempt is made to hide the assessment activities. Asset owner or vendor personnel work with the team to better focus the testing efforts and answer questions about the system. INL has found that facility personnel are often aware of where the problems exist. Involving them in the assessment process can save valuable testing time and ensure that critical or insecure areas are given sufficient attention. The assessment team and vendor or asset owner can teach and learn from each other.

Penetration testing of new systems should be conducted to identify potential impacts that exploitation of vulnerabilities have on ICS functionality before the system is put into production. Testing can also be performed on disconnected development or backup systems to generate representative impacts to ICS functionality.

Associated risk


Penetration testing can pose significant risk to ICS systems. At a minimum, it may slow the networks' response time due to network scanning and vulnerability scanning. Penetration activities may render ICS components inoperable, alter system data, or even cause economic or physical damage by manipulating the physical system. Although this risk can be minimised by the use of experienced penetration testers and rules of engagement, it can never be fully eliminated. A hacker poses the same risk, but there are safer ways of identifying security weaknesses.

A vulnerability assessment can simply identify and report noted vulnerabilities, without putting the system at risk by attempting to exploit them.

Levels of disclosure

Penetration testing simulates a hacker who has targeted the company or specific item of interest. Testers usually have little or no knowledge of the company's network. Security assessment teams are given direct access to the target, with varying levels of information.

The amount of information disclosed to the testers can range from no information to full disclosure of network diagrams, source code, IP addressing information, etc. This is known as black-box versus white-box testing. Any level of information between no knowledge and complete knowledge of the infrastructure to be tested is known as grey-box testing. This concept is illustrated in figure 1 below.




(Figure 1: Black-box versus white-box testing)

ICS owners may request black-box penetration tests with the goal of attaining security certification or meeting regulatory self assessment requirements. However, this ignores the fact that any targeted attack on a system most probably requires some knowledge of the system, and any insider attacker would be in possession of as much information as the system owners. In most cases it is preferable to assume a worst-case scenario and provide the testers with as much information as they require, assuming that any determined attacker would already have acquired this.^c

Grey-box testing is generally the optimal solution because the benefits from both black-box and white-box testing can be leveraged for the particular situation. Figure 2 illustrates this point.

^c www.securitydocs.com/library/3099



(Figure 2: Benefits of black-box and white-box testing combined in grey-box testing)

Focus of testing

A typical cyber security penetration test is focused on the corporate/IT environment and the weaknesses exposed to the outside world that may allow an attacker unauthorised access from the internet. These internet-to-corporate tests are rarely part of an ICS cyber security assessment.

The protocols used in ICSs differ from generic IT protocols. Many ICS vendors use proprietary protocols for inter-process communications. These protocols were developed when ICSs were isolated from the corporate environment and security was not a consideration. Also, the fact that the protocols were proprietary led some vendors to mistakenly believe that an attacker could not exploit them. Communications to field devices often use published industry standard protocols such as Distributed Network Protocol 3.0 and Modbus. These protocols were originally developed to run over serial connections, but were layered on top of TCP/IP for the convenience and efficiency of LAN/WAN communications. Many of these proprietary and industrial protocols lack any means of authentication or integrity checking, and some industry protocols are published with information freely available on the Internet. With ICSs no longer isolated from the corporate/IT world, these insecure protocols put the systems at risk of a cyber attack.

Because of the inherent insecurity in the ICS environment, ICS testing focuses on the security of the ICS electronic perimeter (the communication paths in/out of the ICS network). The team evaluates the network architecture for an appropriate defence-in-depth security strategy, which involves the use of firewalls and the establishment of functional demilitarised zone DMZs. The corporate and ICS networks should not communicate directly, all corporate communications into and out of the ICS network should be brokered through a functional DMZ or other mitigating architecture. Only ICS communications are on the ICS LAN; Internet and e-mail access is not allowed on this network. The team looks for weaknesses in the networks, hosts and applications that could allow unauthorised access into the trusted ICS zone from the corporate or DMZ networks. This includes an evaluation of the placement and configuration of firewalls and intrusion detection devices. Communication links between field equipment and the ICS network are examined for weaknesses. Unlike pentests, which start from the Internet, the ICS

cyber team often begins testing an attack from a corporate client that is sending requests for data to a host inside a functional DMZ or ICS LAN.

Impacts of testing

Typical penetration tests look for known IT vulnerabilities that can be exploited (often with published exploits) to gain unauthorised network access. Penetration testers usually attempt to actually exploit the vulnerabilities to break into the system. The significance of the unauthorised access is determined by the impact on three defined security objectives for information and information systems: confidentiality, integrity and availability (CIA). According to a Federal Information Processing Standard, Publication 199, a loss of confidentiality is the unauthorised disclosure of information, a loss of integrity is the unauthorised modification or destruction of information, and a loss of availability is the disruption of access to or use of information or an information system.^d For typical IT systems, the security goals of CIA are listed in order of importance, with confidentiality considered the most important.

In general, the most significant difference between the ICS and corporate IT domains is the high availability requirement for monitoring and control functionalities, as illustrated in figure 3.

Cyber security is the protection of information transmitted and stored over a computer network. The objectives of cyber security are to:

- Protect confidentiality of private information;
- Ensure availability of information to authorised users on a timely basis (authentication, non-repudiation);
- Protect the integrity of information (i.e., accuracy, reliability and validity).

These objectives can be prioritised differently depending on the physical system under control and the functionality provided by the individual ICS component.

^d U.S. Department Of Commerce, *Standards for Security Categorization of Federal Information and Information Systems*, National Institute of Standards and Technology, Gaithersburg, MD, 2004.



(Figure 3: Generic IT security goals versus ICS security goals)

Vulnerabilities are often exploited during an ICS cyber security assessment. Any exploit development is accomplished on an evaluation or development system and never on an active system. For ICSs, CIA security objectives are in reverse order of priority, with availability considered the most important. Industry personnel may often use the term 'security' to mean availability and reliability. Systems that control the critical infrastructure must constantly operate and the impact of downtime can range from inconvenient to catastrophic. Because public health and safety may be at risk, vulnerabilities found during ICS cyber security assessments at an asset owner's facility are never exploited unless the test can be performed on isolated or offline components. The team works with ICS engineers and other facility personnel to determine the potential impacts the identified vulnerabilities may have on the ICS.

Nothing must be done on the active ICS network that would interfere or disrupt the time-critical operations of the system. In the ICS environment, the CIA security objectives of the IT world are replaced by human health and safety, availability of the system, and timeliness and integrity of the data. This is the major difference between ICS and IT security assessments. This difference also holds true for mitigating strategies. No cyber security solution can be implemented on the ICS network if it interferes with the response of the system. The cyber assessment team must work with industry and vendor personnel to deliver an effective assessment without compromising the safety, availability, or integrity of the ICS.

Testing process overview

This section provides an overview of the testing process focusing on the logistics of an assessment. A description of the overview is provided first to showcase the testing process. The sections that follow detail important topics such as selecting an assessment team, choosing the attack vectors, executing the test and the test report. Also considered is the follow up to a cyber security assessment to see if the problems were resolved or mitigated.

Overview

The process of conducting a cyber security assessment of an ICS is often initiated by a pre-assessment meeting between the leader of the assessment team and key people (network engineers, ICS engineers, instrumentation engineers, security, safety and data users) from the ICS vendor or asset owner's organisation. This meeting will often review the high-level structure of the system and define the system configuration for the assessment (i.e. production, representative laboratory or backup system). This usually includes identifying the key ICS servers and the roles and responsibilities of each on a network diagram. After the ICS structure has been presented, the discussion focuses on identifying initial attack vectors to be included in the test plan. This process is where the attendees openly present their ideas on areas of the ICS that are vulnerable to a cyber attack (attack vectors). Once a list of attack vectors has been generated, each item is ranked according to its potential damage to the ICS if compromised. This ranked list is a starting point for the assessment team's efforts. The pre-assessment meeting also establishes the rules of engagement for the assessment. These rules include declarations of known problems and lists of processes and IP addresses to be excluded during the assessment. Typically, the last item of business for this meeting is to identify the points of contact during testing. This usually means establishing a schedule of people who are on call to the assessment team during the testing period to assist and authorise tests.

Once the assessment begins, the system configuration is fixed in place. The asset owners and operators are prohibited from making changes to the system during testing without coordinating with the assessment team. Often, the ICS administrator (or other suitable nominated individual) will participate in the testing by shadowing the assessment team. This interaction allows the administrator to learn how an attacker would operate inside the ICS. This administrator also acts as an information resource to the assessment team and as the communication medium to the asset owner of current test results. The ICS administrator is essential if the testing is performed on a production system where incremental authorisation is required for each step.

The testing proceeds with the assessment team performing standard tests, such as port and vulnerability scans, to see whether the ICS is susceptible to the current publicly disclosed vulnerabilities. Next, the team starts working on the attack vectors that were identified in the test plan. These items can often be worked in parallel, allowing the cyber lead to divide the team to work on separate efforts. The cyber lead of the assessment team will likely assign a level of effort to each task. These allotments allow the team to cover a larger number of attack vectors so they operate more efficiently. The team does not want to be constrained to searching for a problem in a portion of the system that may be operating securely. The attack vectors specified in the test plan may uncover problems, but may clue the team into other areas that lead to a vulnerability discovery. It is optimal if the assessment team has the freedom to vary from the test plan because they may expose any easily exploitable vulnerabilities an attacker is likely to use.

Depending on the rules of engagement, the assessment team may verbally report a vulnerability as soon as they are able to demonstrate it. Other times, the assessment team may wait until the end of the assessment to informally report on the items found. In either case, the testing will be followed by a detailed written report. It is valuable to the asset owner to


understand the successes as well as the failures because this data provides a measure of how resilient each tested system component is to a cyber attack.

The assessment team may do little to hide their activities on the network. A cyber security assessment of an ICS is not a penetration test in the sense that is common in IT space. A penetration test implies a black-box test where the attackers are working their way inside an organisation starting from the internet. The IT penetration test establishes how far an attacker could penetrate the system. Therefore, the organisation employees are often unaware of the testing to prevent biased results. On the other hand, when asset owners prescribe a cyber security assessment of an ICS, they want to know if vulnerabilities exist inside the hardware and software that make up the ICS and whether the protections (network architecture, functional DMZs, sensors) in place will limit access. It is optimal for the asset owner to have the system administrators and operators work with the assessment team to maximise the testing and facilitate a learning environment for one another. However, even though the assessment team will not try to hide their activities, it is valuable to determine if the network sensors detect the assessment activities. This information assists the organisation in placing alarms and blocks to help detect and prevent a real attack.

Following a cyber security assessment, a number of vulnerabilities are often reported to the asset owner. Once the asset owner or vendor has had a chance to work on these problems, they may request that the assessment team validate the patches. Alternatively, the asset owner may not be able to mitigate a vulnerability because they do not have access to the source code for that application. In this case, the assessment team can assist the asset owner by either working directly with the vendor or through the product users' groups to influence the vendor to fix the vulnerabilities.

The assessment agreement should define roles and responsibilities with respect to disclosure of vulnerabilities identified during the assessment. The ICS owner may require a non-disclosure agreement that prohibits the assessment team from disclosing system and vulnerability information.

The general ICS assessment process overview is summarised in figure 4.



(Figure 4: Assessment process flow chart)

Choosing the assessment team

The asset owner chooses the testing organisation or provider, but may have little control over the actual members of the assessment team. Information about the team members should be provided by the organisation hired to perform the assessment. This information may include certifications, experience, skills and confirmation of background checks.

Certifications are a debated qualification for security testers. At the Defcon 16^e security conference in a talk entitled ‘The pentest is dead, long live the pentest,’^f the presenter suggested that a security certification is about as valuable as ‘a note from your mom.’ It is true that many of the high-profile security researchers do not hold certifications or advanced college/university degrees. Rather, these individuals gained their knowledge by experience. This same population of non-certified researchers is largely responsible for development of the advanced exploitation and defence techniques in which commercial organisations later ‘certify’ students. An individual may hold many certifications and yet have little practical experience. However, while certifications do not guarantee competence, they may provide some measure of the level of training an individual has attained.

The following list highlights some of the available cyber security certifications:

- Certified Information Systems Security Professional (CISSP) - an information security certification accredited by American National Standards Institute (ANSI) International Organisation for Standardization (ISO)/International Electrotechnical Commission (IEC) Standard 17024 and governed by the International Information Systems Security Certification Consortium.
- Global Information Assurance Certification (GIAC) Security Essentials Certification (GSEC), along with other applicable qualifications such as GIAC Certified Incident Handler (GCIH), GIAC Certified Penetration Tester (GPEN) and GIAC Assessing Wireless Networks (GAWN). GIAC is also accredited by ANSI ISO/IEC Standard 17024 and is affiliated to the SysAdmin, Audit, Network, Security (SANS) Institute, a research and education organisation.
- Certified Ethical Hacker - a professional certification provided by the International Council of Electronic Commerce Consultants.
- CHECK approved IT health check service providers approved through the Communications – Electronic Security Group (CESG), utilising the Council of Registered Ethical Security Testers (CREST) and TIGER Scheme Infrastructure Certification Examination.

ICS cyber security assessments differ significantly from standard IT-type assessments. It is imperative that members of the assessment team have experience with assessing ICSs and are aware of the limitations and challenges associated with testing in a production environment. The asset owner should validate the team’s references to ensure that the team has adequate ICS experience. The testing organisation should provide the asset owner with a methodology of how assessments are performed in a production ICS environment. The methodology should include a list of typical tools used by the team and indications of when and how the tools will be used.

The roles and responsibilities of each team member should be clearly defined and communicated to the asset owner. If the assessment is to include a network analysis, at least one team member should have qualified networking experience and possibly network certifications such as those available from Cisco. At least one team member should be familiar with a number of the network protocols unique to ICS (e.g. DNP3, the Modbus suite, PROFINET, PROFIBUS, ICCP, OPC, etc.). This individual is responsible for analysing network traffic and assessing the configurations of network devices such as firewalls, switches and routers. Other team members should be proficient in coding, reverse-engineering, protocol analysis and exploit development. The team members should be familiar with multiple

e. www.defcon.org

f. mirror.sweon.net/defcon16/Speakers/Banks-Carric/defcon-16-banks-carric.pdf

languages such as C, C++, Python, Perl and assembler. Many ICSs have non-Windows operating systems and the asset owner should ensure that team members are familiar with the operating systems used on the target system. If possible, the asset owner should request individuals familiar with the protocols, programming languages, applications and operating systems used by the ICS.

As the assessment team will have access to sensitive information, the asset owner should be provided with confirmation that appropriate background checks have been performed. The control methods for the information acquired during the assessment may be defined in a legal document such as a non-disclosure agreement.

The test plan

It is mutually beneficial for the assessment team and the asset owner to create a test plan before testing begins so that both entities know how the assessment will operate, including the rules of engagement, attack vectors and points of contact. However, the level of effort put into the test plan is a grey area that has to be decided by the asset owner. The asset owner may be more comfortable including a lot of detail in the test plan so that this document can act as a contract with the assessment team. However, the assessment team does not need great detail in this document, the exception being the rules of engagement. In fact, it may be a hindrance to the assessment team for the test plan to include many details (discussed in the next section). Ultimately, the time and money spent creating the test plan may be subtracted from the testing operations budget. This fact could potentially restrict the assessment team from accomplishing some portion of the desired testing.

Selecting the attack vectors

One of the pre-assessment meeting tasks is to establish a set of initial attack vectors to include in the test plan. Many criteria may be used to select these items, but use caution in the level of detail specified. A detailed test plan prescribes exactly what to test, which will ensure that the assessment team covers the items identified. However, all the testing hours might be consumed filling in the details in the test plan without uncovering easily accessible vulnerabilities in other areas not included in the plan. An example might be that the asset owner wants to know if an attacker can take control of the front end processor (FEP) based on the communications allowed from the remote terminal unit (RTU). While this may be a valid concern in this particular installation, it might be overshadowed by the privileges extended to the ICS vendor maintenance connection. The alternative is to specify functions or transitions (discussed below) that present a potential attack vector. An example might be to test whether an attacker can make a network transition from one of the DMZ servers to a server inside the control network.


The most important part of planning a cyber security assessment is that the plan should not constrain the assessment team to approach a problem from only one direction. The test plan should loosely define what to test and never how to test it. This allows the cyber team to use all their skills to accomplish the goals. After all, by definition, a potential attacker is not going to follow the rules of engagement.

Components that provide core ICS functionality should be included in the assessment targets, for example:

- Attack the FEP from the field equipment side (manipulate the RTU or PLC connection);
- Attack the FEP from the ICS network side;
- Attack the application server (e.g. the HMI);
- Attack the real-time database server;
- Attack the historian server.

The descriptions are intentionally vague. The actual attack vector description may include little more than what is listed. These high-level descriptions preserve the flexibility that the assessment team needs to explore the problem in unconventional ways — the way an attacker would operate.

Even though the attack vector descriptions are vague, the goal of each of these items could include additional detail. For example, a common goal may be to demonstrate remote control of a process or server. In the case of a database server, the real question is whether an attacker could manipulate the data stored in the database. By leaving the description vague, the assessment team can attack that server using the database application itself or any other facet of this server such as the operating system or other network processes. In many cases, an attacker can gain control of a server by attacking one process and then leveraging that access to manipulate the true target process (in this case, the database). Cyber security tests structured in this manner tend to expose the easiest way to attack a given server, which is an advantage to the asset owner because the first target found may also be the easiest problem to mitigate.




(Figure 5: Example of supervisory control LAN attack targets)

In addition to the major components of the ICS, other categories make good attack vectors for the test plan. One of these categories is a network transition. An ICS is usually protected behind several layers of network defence from the internet. The goal of a transition would be to gain remote control (by any means) of a server inside the target security zone from a network presence on a lesser security zone. Therefore, many asset owners would like some measure of how far an attacker could penetrate their infrastructure. The test plan may include attack vectors such as:

- Transition from a presence on the corporate LAN to a DMZ server;
- Transition from a presence on the corporate LAN to an ICS server;
- Transition from a DMZ server to an ICS server;
- Transition from the ICCP server to another DMZ server.

The above list is brief, but it could be much longer if an organisation has complex network structures with multiple zones such as: corporate, DMZ, ICS, management, visitor, VPN, field equipment; LANs; and WANs. DMZs are a good example of complex network structures because an organisation may have multiple functional DMZs, including Web, database, historian, application (Citrix farm), Inter-Control Centre Communications Protocol (ICCP) and Oasis.

Another category of tests that make good attack vectors are key functions of the ICS. An example might be data replication. A common ICS configuration is for data to be pushed from the control network to a DMZ server where the data can be polled from hosts on the corporate LAN. Attackers may be unable to get to the ICS itself, but may be able to manipulate one piece of the data replication chain. Therefore, attackers may control all the data to the downstream consumers. If remote HMI functionality is made available on the system, it should be top priority for security assessment activities. Any other remote management functionality, should also be included as attack targets (i.e. vendor VPN access, remote administration of hosts and network equipment, etc.).



(Figure 6: Potential attack vectors between ICS network security zones)

Assessment execution

Once the test plan has been written and the team has been selected, testing begins. Testing is an iterative process of reconnaissance, exploration and exploit development. The reconnaissance phase focuses on identifying targets to attack based on a reason or theory. Exploration is the process of validating whether the target is vulnerable to the attack specified in the theory. Exploit development is the activity that explores the potential of a given vulnerability. Each of these activities is discussed in detail in the following sections.

Reconnaissance

The first part of a cyber security assessment is to identify a target to attack. A number of methods are available to identify such a target. A common practice in a cyber security assessment is actively scanning potential targets. A port scan using tools such as Network Mapper (Nmap)^g quickly identifies the ports on which a host is listening for connections. Because many standard services run on well-known ports,^h the results from Nmap may identify standard services known to have security weaknesses, such as clear-text authentication (i.e. telnet and ftp) or weak authentication. The Nessusⁱ vulnerability scanner performs the same port scanning function as Nmap and, in addition, tries to identify whether the target host is patched against a library of publicly disclosed vulnerabilities. Vulnerability scanners are useful to quickly identify if a host is missing important patches; however, these tools often report false positives. It is, therefore, important to validate every claim that such tools report.

Port scanning tools are as common to a cyber security researcher as a hammer is to a carpenter. However, scanning tools can have drastic effects on some hardware and software. In a perfect world, commercial off-the-shelf hardware and software should be able to stand up to a port scan just like any other request. The truth is that many processes and servers will crash or become unresponsive when the processes and servers are scanned. Unfortunately, this is especially true in the ICS domain. Many field devices currently deployed will crash or become unresponsive from a simple scan. Therefore, cyber security assessments of an ICS have to keep this in mind. For example, these types of tools should never be used on a production system. The crashes themselves may provide useful information. If a process crashes from a scan, this is an indication that the process may be exploited to gain remote access to the host. For systems that require high availability, including critical ICS components, a crash can be a significant vulnerability. Scanning is a preferred method to interrogate a process or server that often reveals targets for attack.

In addition to active host scanning, other passive means, such as monitoring network traffic, may be used to identify targets to attack. Network traffic captures are a great way to identify the protocols and 'big talkers' in the ICS network. Many cyber security assessments begin by creating a span port on each of the key network switches to enable a capture session using tools such as tcpdump.^j Often, cyber researchers will capture traffic on these span ports for a number of hours and then run analysis tools on the data. These analyses produce statistics that identify which hosts are talking to each other, what protocol, and how much data is sent. It is sometimes productive for an attack vector to go after the hosts that are communicating the most. Network captures can also be used to identify clear text communications. The Wireshark^k tool can be used to visually inspect the captured data. This tool has built-in decoders for many popular Internet protocols, which provides easy access to the individual fields in a packet. This ability allows the cyber researcher to look for key fields such as length fields, American

g. nmap.org/

h. www.iana.org/assignments/port-numbers

i. www.nessus.org/nessus/

j. www.tcpdump.org/

k. www.wireshark.org/

Standard Code for Information Interchange (ASCII) strings, null terminators and message identification tags. Identification of any of these fields constitutes a target for attack.

Category	Network reconnaissance tool	Use	Risk
Port scanners	Nmap	Network mapping: Network IP and port detection	Crash / DoS of critical services
Vulnerability scanners	Nessus	Known vulnerability identification	Crash / DoS of critical services
Network monitoring software	Wireshark / TCPDump / Ettercap	Network traffic analysis	Can be performed safely from span port

(Table 1: Popular network reconnaissance tools)

One of the most important places to capture traffic is the electronic perimeter (ingress/egress) points that make up the boundaries between network segments (e.g. boundary between the corporate and the DMZ LANs). In and of itself, mapping the data that traverse these boundaries is not an attack vector. However, this effort will often reveal unexpected transactions that make good attack vectors. This effort is immensely valuable because it can be used to understand the ICS from the network perspective. Very little of this low-level information will be available in the product documentation or from the system integrators. Observing the ICS component interactions will often lead to a functional attack vector. The assessment team should use the traffic capture data as well as the port scan results to build up a master list of port-to-service mappings. This list can be used by the assessment team as a checklist of potential attack vectors to investigate. A list of open ports and associated services should be reported for each ICS component assessed. This list can also be used to validate the accuracy of network diagrams. ICS owners should disable all unnecessary services. See table 2 overleaf as an example.

Port	Service	Security notes	Recommendation summary
21/TCP	FTP	Transports passwords in clear text	Replace with SFTP and restrict access if needed; otherwise remove and block port.
80/TCP	HTTP	Unpatched Web server	Install patches and restrict access if needed; otherwise remove and block port.
XXX/TCP	Proprietary ICS Service	Remotely exploitable by a buffer overflow attack	Restrict access; closely monitor port; work with vendor to obtain patch.

(Table 2: Open ports on host (IP) example)

The ICS domain contains equipment that communicates by means other than IP such as radio, serial and modem lines. Part of the reconnaissance phase should be to trace all the data entry points to the ICS. Sometimes this means physically following cables. While modem and serial lines are not as convenient for an attacker, these represent as real a threat as IP communications.

Cyber security assessments of an ICS should also examine the networking equipment in use at the installation. The configuration files for these devices will identify the access control lists and other deployed protections. For example, the assessment team could compare the firewall rules regulating traffic entering and leaving the control network with the port scan results to see if an attacker could reach the listening services inside the ICS. Other checks include the authentication mechanisms being used by the switches and routers. A common mistake found in ICS networks is that the system administrator will authenticate to a firewall with a secure identification (ID) token over telnet. This transaction is vulnerable to a man-in-the-middle attack because telnet is a clear-text protocol and the attacker can quickly acquire and use the token ID. Network equipment may be ignored on a non-production ICS if it has not been configured for assessment. Disconnected test systems, built to identify vulnerabilities in the ICS software and hardware, may not include network devices with representative configurations. Problems in the networking equipment are likely to be site-specific, which means that this task is more suited to the production system assessment. The main testing goal of a production system assessment may be to identify network security weaknesses that may allow an attacker to reach the ICS, rather than attacking the production ICS.

Network diagrams generated for testing purposes should also be included in the assessment report. Often these diagrams are more accurate than the existing system diagrams.

Exploration

Once a target has been identified, the assessment team attacks the system. Cyber security attacks can be summarised as 'exploiting assumptions'. The reason for this definition is that many of the vulnerabilities found in software and hardware result from using a function in a manner that the designer did not anticipate.

In the ICS domain, the assessment team may not understand all the components of the system. This understanding is not required for many cyber security tests. To be clear, it is true that an attacker would need some understanding of the particular ICS for a surgical attack - to force point X to value Y while showing the operator value Z - but it may not be that hard for the attacker to perform only half of the equation. An ICS is really a collection of servers and processes. When an ICS is thought of in these terms, the attack process is somewhat generalised.

The attacker begins the intrusion by conducting some documentation research. Many ICSs are deployed with elaborate help systems that include default configuration settings. It is common for the ICS documentation to include default account and password information. In fact, some legacy systems prohibit the end user from changing default passwords. Alternatively, the attacker could turn to the internet because the system vendor or user group may have posted manuals or system information online. In this stage of a cyber assessment, the researcher is probably only interested in a general idea of what this process or component supplies. The attacker is not interested in all the features of this component. In many cases, the attacker will not bother to look at the documentation of a component until he has already exploited it or is very close and is missing some small detail of information.

Attacking a network process requires the attacker to obtain network communication captures of the normal operations of the target application. Sometimes, the attacker will capture the traffic from the process startup. Startup traffic is unique because it will likely demonstrate a connection scenario. The attacker uses this information to create a client with which to create a new connection to the target process. Alternatively, the attacker could start a man-in-the-middle attack and redirect the in-progress network traffic stream to and from the target so that it first passes through the attacker computer. In either case, the attacker will likely manipulate a protocol field that was observed in the earlier network captures. For example, if the protocol is sending a length field followed by a quantity of data, the attacker can simply modify the length field in the network packet. The intended receiver of these data (the victim) now reads too much or too little data from the network, which causes problems in the parsing state machine. This is an example of the attacker manipulating the designer's assumption that the length field in the packet accurately reflects the amount of data to follow. The result in the victim process is that a receive buffer now contains too much or too little data, which may lead to memory overwrites and process crashes. The attacker will attach a debugger to the victim process before the start of the attack in case a crash occurs. The attacker will then analyse the traffic processing and develop an exploit for this vulnerability.

In the previous example, the attacker had already identified a specific field in the data stream to manipulate. Data sent in clear text is not necessarily in plaintext, which means that it may be difficult for the attacker to identify key fields in a large stream of binary data. In this case, the attacker may choose to fuzz the data stream sent to the victim. Protocol fuzzing, or just fuzzing, is the process of sending semi-valid data to a process and observing its behaviour; length fields are set to extremes and other boundaries are stressed. This method may expose vulnerabilities in a process even when the attacker knows little about the protocol. This method works because in addition to length fields, many common network fields are modified to cause disruptions in the victim process. Examples of these fields include ASCII NUL character terminators at the end of a string, message identification tags and data type specifiers.

Instead of, or in addition to, examining the network traffic to learn of a protocol, the attacker could examine the binary. Tools such as IDA Pro¹ allow a researcher to reverse engineer a binary from machine code to assembly instructions. A skilled researcher can use the assembly to decipher how a process works and sometimes recover the original source code. In this manner, the researcher looks for instances of programmer mistakes and shortcuts, which are potential vulnerabilities. The researcher crafts a message to send the process and redirect the binary logic to the vulnerable code.

In some ICS assessments, the team may be given source code for an ICS process because either the asset owner developed the application or because the vendor is participating with the assessment. In this case, the researcher searches directly for problems in the code. The number of publicly disclosed vulnerabilities should be evidence that attackers do not require source code to find vulnerabilities. However, having the source code available for inspection speeds up the vulnerability identification process.

1. www.hex-rays.com/idadpro/

Asset owners may take the black-box approach where they provide little detail to the assessment team in an attempt to see what an attacker could do without any inside knowledge. Assessments where the assessment team is given no source code and little other information identify fewer vulnerabilities. On the other hand, if the assessment team is given source code and/or information about the installation, they can perform a more in-depth assessment, where they potentially drill deeper into each process.

Another area that is gaining popularity in the ICS domain is Web and database applications. These applications are commonly used to allow corporate users to view data from the ICS. The assessment team may find additional attack vectors by examining these applications for problems such as Structured Query Language (SQL) injection^m or Cross-Site Scriptingⁿ (XSS) problems. Attackers can use these problems to make a network transition from the corporate LAN to a DMZ server or even to an ICS server.

In addition to the areas listed above, the assessment team may check a number of other items as they look for attack vectors. The items in the following list have been reported in a number of ICS security assessments. The high level security weaknesses should be included in the assessment plan, if applicable:

Published vulnerabilities:

- Use of vulnerable remote display protocols;
- Secure Shell daemons that allow older versions of the protocol and are vulnerable to a downgrade attack;
- Anti-virus and spyware programs that do not have current signatures or are updated in such a manner that open an attack vector;
- Lack of a patching process/schedule leaves the ICS hosts open to attack from publicly disclosed vulnerabilities;
- Domain hosts using or storing antiquated LanMan hashes, which can be cracked using a dictionary attack;
- Backup software vulnerabilities that allow the attacker to manipulate data or server.

Web vulnerabilities:

- Web HMI vulnerabilities;
- Secure Sockets Layer man-in-the-middle attacks where the attacker takes advantage of self signed HyperText Transfer Protocol over Secure Socket Layer (HTTPS) certificates.

Input validation vulnerabilities:

- Buffer overflows in ICS services;
- SQL injection.

Improper authentication:

- Authentication bypass, e.g. client-side authentication;
- Use of standard IT protocols with clear-text authentication;

m. en.wikipedia.org/wiki/SQL_injection

n. en.wikipedia.org/wiki/Cross_site_scripting

- Unprotected transport of ICS application credentials.

Improper access controls (authorisation):

- Wireless LAN access that can be used to get to the control network;
- Blank system administrator^o password on a Microsoft SQL Server database, which allows remote administrator access to the database and the server itself;
- VPN configuration problems that unintentionally allow clients unfettered access to the corporate, DMZ, or control LAN;
- System management software that allows central management of multiple servers may allow an attacker easy access to the same hosts;
- Common processes (any process that is installed and listening on multiple boxes), which if compromised, provide access to multiple hosts;
- Weak firewall rules;
- Circumvented firewalls;
- Shared printers that span security zones. This may provide a network transition that does not traverse the firewall;
- Unsecure network device management.

ICS data and command message manipulation and injection;

Database vulnerabilities;

Unnecessary or risky services and applications:

- Internet/e-mail access from within secure zones (DMZ, SCADA) may allow malware inside these protected zones.

Poor network monitoring.

Exploit development

Once a problem has been identified, the assessment team may optionally develop an exploit for the vulnerability. In an ICS cyber security assessment, exploits are created for several reasons. First, the asset owner and operator may want proof that the problem can be exploited on their system. An exploited vulnerability dispels any doubt that the vulnerability is real. Second, not all vulnerabilities are exploitable. Therefore, the team may attempt an exploit to provide accurate mitigation recommendations.

There are also reasons not to develop an exploit. Exploit development can take significantly more time than the additional value it adds to the assessment. For example, if an assessment team finds multiple vulnerabilities in an ICS, the team may choose to develop an exploit for only the first few problems. The reason is that they have already earned the trust of the asset owner and now they are trying to cover as much of the ICS territory as possible. In this case, it is sufficient for the assessment team to demonstrate the crash without an exploit.

^o www.ca.com/us/securityadvisor/vulninfo/vuln.aspx?id=5705

Likewise, the assessment team is unlikely to chain exploits together because this effort consumes testing time and generally does not uncover additional vulnerabilities. For example, if the team has demonstrated an exploit that enables a network transition from the corporate LAN to a DMZ server as well as another exploit for a server in the same DMZ, it can be inferred that the two exploits could have been chained together. The researchers could develop a more complex exploit to perform both operations together, but this task is a trade-off to additional vulnerability identification. The benefit from a vulnerability assessment is proportional to the number of vulnerabilities that are identified for remediation.

Assessment reporting

Because the primary product of a cyber security assessment is the report, this section presents suggestions on how to maximise the value from the assessment report. Ideally, the cyber report should be able to meet the needs of many different audiences. The report needs to have high-level language appropriate for managers, as well as detailed technical information for the engineer responsible for mitigating the reported vulnerabilities. Figure 7 sets out possible headings for this report.

This outline is flexible enough to allow the assessment team to report on all the attack vectors regardless of whether the team accomplishes the goal of the attack vector. For example, the team may work on several attack vectors ('targets' in the report vernacular), each of which could uncover zero to many vulnerabilities. Any given vulnerability may be exploitable by zero to many exploits that could each yield different results.

The asset owner should identify the types of information and levels of detail and formality desired in the report, because reporting has to be funded out of the assessment budget.


Report executive summary

The executive summary is a less technical summary of the test results. This section lists the vulnerabilities uncovered as well as provides some measure of the effort required to mitigate these problems. Such vulnerabilities or other security issues should be clearly ranked, to enable the organisation to prioritise its remediation efforts. The executive summary also lists some of the positive items observed or tested so management gains a balanced evaluation of the tested ICS system.

Report introduction

The introduction section provides a high-level description of the cyber security assessment including the what, why and rules for this assessment. The objective provides a general description of why this assessment is being performed. The significance description is a description of why this particular system is being tested. The rules of engagement section lists the rules, controls and limitations under which the assessment is conducted.

Cyber Report Template Document Flow Diagram



(Figure 7: Sample cyber security assessment report outline)

The following template may be included:

This report describes the cyber security assessment of the <Utility><backup> <Vendor> system and electronic perimeter conducted <range of dates>. The <backup> system assessed included core components of the operational configuration, providing the assessment team with representative data while effectively isolating the assessment operations from affecting normal utility operations. In addition, the assessment team carefully followed the pre-determined rules of engagement.

- *Rules of engagement*

The rules of engagement describe the constraints within which the assessment team's assessment activities are to be performed. These rules help ensure the safety of all personnel involved in the assessment, the security of sensitive information used in or generated by the assessment process and the integrity of the production environment during the engagement.

- *Safety requirements*

This section can include the safety requirements identified by the asset owner, including procedures and Personal Protective Equipment (PPE) required to operate at their facility.

- *Facility access requirements*

This section can include the facility access requirements identified by the asset owner for access to their facilities.

- *Computer security requirements*

This section can be used to document any waivers that were required and any special considerations for using facility resources.

- *Boundaries*

Set the specific limits necessary to preserve the integrity of the production environment during the assessment. (Table identifies devices that are approved for assessment team personnel to access and those for which access is prohibited. This table should come from the assessment plan.

Inclusions	Description of included devices	Exclusions
<123.456.78.xx>	<Backup server for late night tables.>	<123.456.78.9>

(Table 3: Example of assessment boundary list)

Report target(s)

The target sections are where the team reports on the individual attack vectors they worked. These sections allow the researcher to explain each attack vector with pictures and examples for clarity. For example, each target section may begin by describing the component being tested and the goal of this test. This introduction is followed by dialogue that helps the reader understand how the researcher performed the testing of this attack vector. The dialogue format

also allows the researcher to explain why each step was significant. It is important in a cyber security assessment to capture the failures as well as the successes of the assessment team. Failure to achieve an attack vector goal may be an indicator of good security practices.

If the asset owner would rather have a high-level final report, the team could create individual vulnerability reports for each attack vector. In this case, the master report would only include brief descriptions of each vulnerability and the scoring metrics. The reader would be free to open the specific detailed report if desired. The data should be the same for either report style; it is just a matter of preference where the details are located.

- *Target significance*

The criticality of the targeted component, security implications and other factors that contributed to the target's prioritisation can be documented in this section. This information should also be documented in the assessment plan.

- *Rules of engagement*

Target-specific rules of engagement should be listed in the assessment plan and adherence or deviations can be documented in the assessment report.

- *Vulnerabilities*

If a vulnerability is uncovered, this section allows researchers to describe in detail what was found. This description includes all the required information (screen shots, code snippets, etc.) that someone would need to reproduce the test conditions, such as the asset owner, or more likely, the vendor, who should be informed so as to enable subsequent securing of the product(s) involved.

- *Exploit(s)*

Once a vulnerability has been identified, the next task is to determine the impact of this vulnerability. One of the main reasons that the assessment team will create one or more exploits for a vulnerability is to understand the impact of the problem. Exploit development takes time, but is often worth the cost because the result is a definitive indicator of whether the problem is an issue to the asset owner. Some vulnerabilities produce a situation where an attacker can execute arbitrary code on the victim server while others allow privilege escalation or a denial-of-service condition. Information gathered while attempting an exploit help to score and categorise a vulnerability. This section should include some discussion of the level of difficulty required to create the exploit.

- *Metrics*


Metrics are an important part of vulnerability assessments. Metrics provide a methodology for evaluating the risk associated with a vulnerability. The Common Vulnerability Scoring System (CVSS) is a standardised method of scoring vulnerabilities in a way that represents the risk to an individual organisation's unique environment.^p The CVSS v2 scoring method is a cyber security industry standard which allows vulnerabilities to be prioritised according to the actual risk they pose to the organisation.

Whilst this guide further considers the use of CVSS below, asset owners should be aware that there are other vulnerability assessment methodologies available.

^p Mell, Peter, Karen Scarfone, and Sasha Romanosky, *A Complete Guide to the Common Vulnerability Scoring System Version 2.0*, FIRST Forum of Incident Response and Security Teams, June 2007.

CVSS is a free and open standard. A CVSS v2 scoring guide (www.first.org/cvss/cvss-guide.html) and calculator (nvd.nist.gov/cvss.cfm?calculator&version=2) are available.

CVSS is composed of three metric groups: base, temporal and environmental, each consisting of a set of metrics, as shown in figure 8.



(Figure 8: CVSS Metric groups)

These metric groups are described as follows:

Base: represents the intrinsic and fundamental characteristics of a vulnerability that are constant over time and user environments.

Temporal: represents the characteristics of a vulnerability that change over time but not among user environments.

Environmental: represents the characteristics of a vulnerability that are relevant and unique to a particular user's environment.

The temporal and environmental groups allow ICS owners to incorporate contextual information that more accurately reflects the risk to their unique environment. This allows them to make more informed decisions when trying to mitigate risks posed by the vulnerabilities.

- *CVSS v2 Base Metrics*

The base metric group captures the characteristics of a vulnerability that are constant with time and across user environments. The access vector, access complexity, and authentication metrics capture how the vulnerability is accessed and whether or not extra conditions are required to exploit it. The three impact metrics measure how a vulnerability, if exploited, will directly affect an Information technology (IT) asset, where the impacts are independently defined as the degree of loss of confidentiality, integrity, and availability. CVSS v2 base scoring metrics are summarised in table 4. To aid understanding, yellow designates lower risk and red represents higher risk metric values.

Base metrics	Metric value	Metric description
Access vector	Local	Requires the attacker to have either physical access to the vulnerable system or a local (shell) account.
	Adjacent network	Requires the attacker to have access to either the broadcast or collision domain of the vulnerable software, local IP subnet, for example.
	Network	The vulnerable software is bound to the network stack and the attacker does not require local network access or local access, aka 'remotely exploitable.'
Access complexity	High	Specialised access conditions exist.
	Medium	The access conditions are somewhat specialised.
	Low	Specialised access conditions or extenuating circumstances do not exist.
Authentication	Multiple	Exploiting the vulnerability requires that the attacker authenticate two or more times, even if the same credentials are used each time.
	Single	The vulnerability requires an attacker to be logged into the system (such as at a command line or via a desktop session or Web interface).
	None	Authentication is not required to exploit the vulnerability.
Confidentiality impact	None	There is no impact to the confidentiality of the system.
	Partial	There is considerable informational disclosure.
	Complete	There is total information disclosure, resulting in all system files being revealed.
Integrity impact	None	There is no impact to the integrity of the system.
	Partial	Modification of some system files or information is possible, but the attacker does not have control over what can be modified, or the scope of what the attacker can affect is limited.
	Complete	There is a total compromise of system integrity. Complete loss of system protection, resulting in the entire system being compromised.
Availability impact	None	There is no impact to the availability of the system.
	Partial	There is reduced performance or interruptions in resource availability. An example is a network-based flood attack that permits a limited number of successful connections to an Internet service.
	Complete	There is a total shutdown of the affected resource. The attacker can render the resource completely unavailable.

(Table 4: CVSS v2 base scoring metrics)

- *CVSS v2 Temporal Metrics*

The temporal exploitability metric measures the current state of exploit techniques or code availability. Public availability of easy-to-use exploit code increases the number of potential attackers by including those who are unskilled, thereby increasing the severity of the vulnerability. The effectiveness of available work-around mitigations is used to adjust the temporal score.

Temporal metrics	Metric value	Metric description
Exploitability	Unproven	No exploit code is available, or an exploit is entirely theoretical.
	Proof-of-Concept	Proof-of-concept exploit code or an attack demonstration that is not practical for most systems is available. The code or technique is not functional in all situations and may require substantial modification by a skilled attacker.
	Functional	Functional exploit code is available. The code works in most situations where the vulnerability exists.
	High	Either the vulnerability is exploitable by functional mobile autonomous code, or no exploit is required (manual trigger) and details are widely available. The code works in every situation, or is actively being delivered via a mobile autonomous agent (such as a worm or virus).
	Not Defined	Assigning this value to the metric will not influence the score. It is a signal to the equation to skip this metric.
Remediation level	Official Fix	A complete vendor solution is available. Either the vendor has issued an official patch, or an upgrade is available.
	Temporary Fix	There is an official but temporary fix available. This includes instances where the vendor issues a temporary hotfix, tool, or workaround.
	Workaround	There is an unofficial, non-vendor solution available. In some cases, users of the affected technology will create a patch of their own or provide steps to work around or otherwise mitigate the vulnerability.
	Unavailable	There is either no solution available or it is impossible to apply.
	Not defined	Assigning this value to the metric will not influence the score. It is a signal to the equation to skip this metric.
Report confidence	Unconfirmed	There is a single unconfirmed source or possibly multiple conflicting reports. There is little confidence in the validity of the reports. An example is a rumour that surfaces from the hacker underground.
	Uncorroborated	There are multiple non-official sources, possibly including independent security companies or research organisations. At this point there may be conflicting technical details or some other lingering ambiguity.
	Confirmed	The vulnerability has been acknowledged by the vendor or author of the affected technology. The vulnerability may also be Confirmed when its existence is confirmed from an external event such as publication of functional or proof-of-concept exploit code or widespread exploitation.
	Not Defined	Assigning this value to the metric will not influence the score. It is a signal to the equation to skip this metric.

(Table 5: CVSS v2 temporal scoring metrics)

- *CVSS v2 Environmental Metrics*

Different environments can have an immense bearing on the risk that a vulnerability poses to an organisation and its stakeholders. The CVSS v2 environmental metric group captures the characteristics of a vulnerability that are associated with a specific environment. For this report, generic ICS security requirements are used to score generic ICS vulnerabilities.

Environmental metrics	Metric value	Metric description
Collateral damage potential	None	There is no potential for loss of life, physical assets, productivity or revenue.
	Low	A successful exploit of this vulnerability may result in slight physical or property damage. Or, there may be a slight loss of revenue or productivity to the organisation.
	Low-medium	A successful exploit of this vulnerability may result in moderate physical or property damage. Or, there may be a moderate loss of revenue or productivity to the organisation.
	Medium-high	A successful exploit of this vulnerability may result in significant physical or property damage or loss. Or, there may be a significant loss of revenue or productivity.
	High	A successful exploit of this vulnerability may result in catastrophic physical or property damage and loss. Or, there may be a catastrophic loss of revenue or productivity.
	Not defined	Assigning this value to the metric will not influence the score. It is a signal to the equation to skip this metric.
Target distribution	None	No target systems exist, or targets are so highly specialised that they only exist in a laboratory setting. Effectively 0% of the environment is at risk.
	Low	Targets exist inside the environment, but on a small scale. Between 1% and 25% of the total environment is at risk.
	Medium	Targets exist inside the environment, but on a medium scale. Between 26% and 75% of the total environment is at risk.
	High	Targets exist inside the environment on a considerable scale. Between 76% and 100% of the total environment is considered at risk.
	Not defined	Assigning this value to the metric will not influence the score. It is a signal to the equation to skip this metric.
Security requirements	Low	Loss of [confidentiality / integrity / availability] is likely to have only a limited adverse effect on the organisation or individuals associated with the organization (e.g., employees, customers).
	Medium	Loss of [confidentiality integrity availability] is likely to have a serious adverse effect on the organisation or individuals associated with the organisation (e.g., employees, customers).
	High	Loss of [confidentiality / integrity / availability] is likely to have a catastrophic adverse effect on the organisation or individuals associated with the organisation (e.g., employees, customers).
	Not defined	Assigning this value to the metric will not influence the score. It is a signal to the equation to skip this metric.

(Table 6: CVSS v2 environmental scoring metrics)

Security requirements metrics enable ICS owners to customise the CVSS v2 score depending on the importance of the affected component to their own organisation, measured in terms of confidentiality, integrity, and availability. DoS vulnerabilities in ICS components that require high availability will receive higher criticality scores than they otherwise would. The effectiveness of available work-around mitigations is used to adjust the temporal score. CVSS v2 environmental scoring metrics are summarised in table 6:

- *NVD vulnerability severity ratings*

The US National Vulnerability Database (NVD) supports CVSS v2 scoring and maps base scores to vulnerability severity ratings.

NVD provides severity rankings of ‘Low’, ‘Medium’, and ‘High’ in addition to the numeric CVSS scores, but these qualitative rankings are simply mapped from the numeric CVSS scores:

CVSS score	NVD severity rating
0.0 – 3.9	Low
4.0 – 6.9	Medium
7.0 – 10.0	High

(Table 7: Vulnerability severity ratings)

These severity mappings can be used for the other CVSS scores as well. Table 7 shows the mapping between numeric CVSS scores and qualitative rankings. Vulnerabilities with high CVSS scores represent a higher risk to the organisation and should be prioritised for remediation ahead of vulnerabilities with lower scores.

- *Vulnerability Mitigation(s)*

Identifying a vulnerability is only half the battle; the real value to the asset owner is for the assessment team to provide applicable feedback to help mitigate the uncovered problems. Multiple methods may be available to mitigate a given vulnerability, which means that the report should list the most appropriate solutions in order of preference. If the root cause of the problem cannot be addressed, the team should provide guidance on other possible options to be used in the meantime. The mitigation descriptions should include sufficient detail that the person who is assigned to fix the problems will not have to repeat any of the assessment team’s efforts in order to understand the vulnerability. The report should include any assumptions in the mitigation recommendations so the asset owner is able put these suggestions into context. An assessment team recommendation may not be possible because the team does not understand the organisation’s architecture, company policies, etc.

Attack scenarios

Once the team finishes their assessment activities, they will have a good idea how the vulnerabilities they found could be used by an attacker. Therefore, it is often beneficial for the team to list scenarios to assist the asset owner in visualising the overall impact of the current problems. One possible method to present this information is for the team to present scenarios in three categories: demonstrated, probable and worst case. Breaking the scenarios into these

categories makes it easier to see the assumptions that are built into the scenarios. Well-developed scenarios will help the asset owners prioritise their mitigation efforts to secure the ICS.

Network assessment

If the ICS owner requested a network evaluation as part of the ICS cyber security assessment, the results can be presented in a separate section. An assessment of the production networking equipment is of highest value to ICS owners who can do little to remediate the actual vulnerabilities in the ICS. Network level defences are almost the only option in protecting the ICS. In-house or component tests may include only a minimum set of hardware that is required to allow the ICS to function. In this case, the team may suggest ways to use networking equipment in the field to implement monitoring and protection features.

Report conclusion

The report conclusion is used to recap the vulnerabilities that were found and identify the likelihood of mitigation. The asset owners use this summary as a scorecard for the cyber security assessment of their ICS. As in the Executive Summary, vulnerabilities and identified security issues should be clearly ranked, to enable the organisation to prioritise its remediation efforts. The conclusion also is a good place for the team to state observations or recommendations that did not otherwise have a logical home in the report.

Vulnerability mitigation and vendor engagement

When the assessment team has finished its work and submitted the assessment report, the asset owner is potentially left with a list of problems on which to work. The ICS administrator may be unable to mitigate all the problems found. For example, imagine the team finds a problem in the FEP communication protocol such that this server can be exploited from the field equipment side of the network. The asset owner is unlikely to have the source code for the FEP software to just fix the bug and recompile. The ICS administrator cannot block the data stream with a firewall rule because this would break the ICS; the FEP could no longer communicate with the field equipment. The asset owner is thus dependent on the ICS vendor to fix this problem.

There have been many different vendor responses to cyber security assessments of ICS hardware and software. Some vendors have their own security programs and are actively seeking security audits from outside sources while others just now are starting to consider security an important part of their product line. The ICS community has long equated availability with security. This definition has evolved as the threat from cyber attack has become increasingly realistic. Vendors do not want security vulnerabilities in their products any more than the users and asset owners do. Yet, a vendor may be slow to fix a vulnerability because of the level of effort required or for other (maybe political) reasons.

If the vendor was involved with the cyber security assessment or is currently engaged with the asset owner through a service contract, the vendor may issue an early vulnerability patch that needs to be validated. Asset owners may be able to perform this testing with internal personnel, or they may need to bring back the assessment team for this work. Asset owners should consider this possibility when they procure a cyber security assessment of their ICS. Alternatively, if the vendor does not issue a speedy patch, the asset owner may request a follow-on report from the vendor that details what the vendor intends to fix and when.

Users and owners of ICSs will continue to find vulnerabilities that they are dependent on the vendor to fix. One method to influence ICS vendors to make changes in their product is by

interaction with the vendor user's groups. Many ICS vendors hold frequent user group meetings where they interface with the owners and operators of their products. These forums have been used to educate users on security issues and to rally leveraging support for the vendor to make changes. Also, these meetings are a good place to collaborate security testing plans. For example, perhaps the ICS vendor has already performed a cyber security assessment of the product, yet the users would like to see additional testing. The user group setting is an appropriate place to form a consortium to share the load and cost of additional cyber security testing.

When all else fails, the asset owner may decide to publicly disclose a vulnerability. The full disclosure method is and has been controversial since locksmiths introduced it in the 19th century. Asset owners do not want to have the vulnerabilities in their ICS made public any more than the vendor.

Assessment variables

This section covers variables that affect how an assessment is performed. For example, different variables exist for testing performed on a production system than one performed in a laboratory setting. The following list describes many of these variables and explains some of the implications on a cyber security assessment of an ICS.

The assessment budget

The assessment budget will always affect the amount of testing that is performed. The budget usually specifies a number of hours allowed for the testing and ICS assessments generally begin with a set number of attack vectors. The testing hours are divided among those attack vectors. Cyber security testing is research, which implies that the end goal is unknown at the start. Therefore, one attack vector may take longer than anticipated. By allotting a given number of testing hours to each attack vector, the team can re-evaluate (based on the momentum of the test) whether to continue with an attack vector when the allocated hours are expended. More hours enable deeper evaluation of each attack vector or the potential for more attack vectors to be attempted.

Prior Information

The assessment team is given information about the ICS before the testing begins. For example, the assessment team is provided with network diagrams and firewall rules prior to the start of the hands-on work. This information helps the assessment team optimise the strategy in dividing the team labour and attempting the attack vectors. Alternatively, the assessment team spends the initial test time absorbing this same data.

Source Code

If the assessment team has access to the ICS source code, they perform a more in-depth test because they can drill inside the processes much faster than if each binary must be reverse engineered. Access to source code enables the assessment team to specify exactly (file and line number) where a vulnerability takes place. If the team does not have the source code to a binary they are attacking, the problems they find can only be expressed in terms of function. The assessment team will not be able to identify which source file produced the problem code they are referencing, especially considering that compilers optimise code. In this case, the team will perform more of a breadth test because they will likely use more blind testing methods such as fuzzing.

Laboratory assessments

Laboratory assessments allow the team to perform a more in-depth test because the team has a number of freedoms which allow them to go into more detail than if they were working onsite. For example, tasks such as reverse engineering protocols/binaries, fuzzing processes, source code review and exploit development all take time. A laboratory assessment is more likely to have the time to assign a researcher to work on a single attack vector for several weeks as opposed to an onsite assessment where the entire assessment may span only a few weeks. Another significant advantage of a laboratory assessment is the ICS will be separate from the production version. This fact means the team will have a green light to non-destructively test any and all parts of the ICS without the possibility of causing real world impact.

Onsite assessment

There are several limitations for onsite assessments which imply that this test will be more of a breadth-type assessment⁹. An onsite assessment requires special consideration because the team will be working closely with a production ICS. This situation warrants that the team will have limited access to the ICS and should work with much more caution and oversight. The increased oversight for onsite assessments adds considerable time to each test, which limits how far the assessment budget can be stretched. Many of the activities performed in a laboratory should not be attempted onsite. For example, it is unacceptable to fuzz a data stream for a production ICS because the goal of fuzzing is to cause a crash. Simple activities, such as port scans may have detrimental effects on a production ICS. The assessment team's activities in this case should be mostly passive. Onsite assessments do have the advantage of examining the real system. Even the best-planned laboratory assessment never fully mimics the conditions of the production system. Therefore an onsite assessment is effective when the goal of the assessment is to validate results and theories created in the laboratory.

Rules of engagement

The rules for a cyber security assessment will directly influence the test results. If the rules specify that the assessment team does not touch a range of IP addresses, these hosts will not be tested. However, there are more subtle results of the rules of engagement such as when the rules prescribe a test should originate from a given point of presence on the network. The firewall may restrict traffic from this starting point, which limits access to a vulnerable service or host. If there are other ways to reach the vulnerability in production installations, this rule has simply masked the problem for the duration of the assessment. It also is common in ICS cyber security assessments for the ICS administrator to exclude items that are identified as restricted during testing. Meanwhile, those items may exist in the field for months to years (typical lifespan of an ICS is 15 years) before the entire infrastructure can be upgraded. Because these items were excluded from the assessment, the full implications of potential vulnerabilities are not understood.

Vendor involvement

ICS software is the major source of risk and obstacles in securing ICSs. Vendor involvement is necessary for remediating ICS product vulnerabilities and requirements that prevent the application of security concepts.

If vendors are involved with the assessment, they may be willing to share information about the system that the assessment team would otherwise have to learn on their own. However, this information and the results of the assessment are likely to be bound by legal agreements that limit with whom the data is shared. These decisions should be agreed on during the pre-assessment meeting, or prior to the start of testing.

⁹ A possible exception is if an assessment takes place whilst the site or system is offline, for maintenance or is not yet fully commissioned.


Alternative methodologies

The actual testing performed on an ICS takes many forms. Until now, this document has focused on attacking an ICS or actively looking for vulnerabilities in hardware and software. This section compares and contrasts this method with several other methods.

There is never enough money to test every part of an ICS. Therefore, the funding organisation plays a large role in determining which of the following tests should be completed under a given assessment budget. This section identifies testing possibilities as well as provides information that will help the asset owner in the assessment plan decision process.

Laboratory assessment

A laboratory assessment is one in which the ICS is offline from the production system. This replicate system should be functionally as close to the production system as possible so the testing mimics the production conditions. Many asset owners have development or test facilities which may be largely pre-configured to the ICS under consideration. Laboratory assessments are often composed of a minimal set of equipment such as in figure 9.



(Figure 9: Sample ICS configuration for a laboratory assessment)

Pros: Because this type of assessment is offline from the production system, the assessment team usually has approval to non-destructively test any part of the system. The assessment team need not worry about causing harm to production processes or people and is more likely to have the time required to perform in-depth tests of the ICS.

Cons: Laboratory systems never truly represent production systems. For instance, these systems are not operating under a production load. The system illustrated in figure 9 is a typical laboratory ICS configuration and does not include all the special connections (VPN, radio links,

vendor connections, etc.) that exist on a real system. Also, a laboratory assessment will not be able to simulate the operators and how they will react to events on their system.

Resources and prerequisites for testing

The minimal system shown in figure 9 still represents a significant amount of extra hardware and software that the asset owner must set aside for the duration of the assessment, though this will not be the case where the asset owner has a pre-existing facility.

The cost of a laboratory assessment is low if an existing back-up, development, test, or new (pre-deployment) system is available. The cost is high if the system must be procured and set up for testing (but it can then be used for future assessments).

Potential impact of the testing

The test equipment may be damaged during testing. Many of the tasks performed during a cyber security assessment of an ICS could result in one or more of the databases containing invalid data. If this system is actually a backup or standby system, it may take time to restore the system to a valid state before it is placed back in service.

Description of the types of results to expect

A laboratory assessment is likely to uncover vulnerabilities in the hardware and software that make up the ICS because the team will have the freedom to test many attack vectors that are not tested on a production system.

Level of security assurance to expect

A laboratory assessment will provide a good measure of how resilient the ICS is to attack. The focus of this assessment is the ICS, not the layers of defence in front of it. This assessment should be able to identify what protections are currently in place and where additional measures could be added at the ICS level.

Example

An example of a laboratory assessment is one in which the asset owner provides a set of equipment and software that mimics the typical deployed system. This test equipment also is a backup or standby system. The assessment is followed by a detailed report that lists vulnerabilities and suggested mitigations. This report will include the team's observations (good and bad) as well as configuration and architecture suggestions.

Conclusion

A laboratory assessment is most effective when the goal is to search for vulnerabilities within the processes and protocols that implement the ICS. This may not be of much value if the owner cannot mitigate the identified vulnerabilities.

INL assessments on ICS vendor products have found vulnerabilities that cannot be mitigated using defence-in-depth strategies because the ICS design requires them to be exposed to less secure zones. Experience has been that some ICS vendors do not address all vulnerabilities provided to them.

Ideally, the ICS vendor is committed to security and quickly patches reported vulnerabilities. An asset owner can help ensure that the money spent on assessing vendor products does not go to waste by creating an enforceable agreement before testing starts. For example, an owner could ask the vendor to prove its commitment by signing an agreement that if the identified

vulnerabilities are not patched within 90 days of notice, those vulnerabilities will be shared with other users of the ICS product (i.e. presented at the next users' group meeting).

Decision Criteria

The following questions can be answered to determine whether a lab assessment should be performed. The associated decision graph is illustrated in figure 10.


Do I have a backup, test, development, etc. system? Or, can I purchase one?

No - Do not perform lab assessment.

Yes - Does my ICS vendor respond to vulnerabilities that are provided to it?

No - Do not perform a laboratory assessment, or, place little emphasis on this process. (Findings will give an indication of the ICS security posture and identify vulnerabilities to monitor for exploitation.)

Yes - Perform lab assessment.



(Figure 10: Lab assessment decision graph)

Production system

A cyber security assessment of a production ICS is performed at the asset owner's location while the ICS is in production. This means that every ICS feature normally used will be present and active for the test.

Pros: Testing the actual production system provides a true evaluation of the real world. None of the assumptions included in laboratory assessments apply because the team is working with a live situation. The impact of an exploit will be physically evident. Therefore, the team does not have to extrapolate their findings to match the 'real' system.

Cons: The risks associated with performing a cyber security ICS assessment on a production system are considerable. For example, the assessment team's activities could cause the ICS to crash or enter unknown states. These disruptions could be costly and even harm property and persons. The team will, therefore, be granted limited access to the ICS and the testing will proceed slowly and cautiously, because there is a large amount of precaution and oversight for this situation. The team should obtain specific permission for each testing activity before it is initiated.

Resources and prerequisites for testing

The asset owner will need to provide the team with access to network span ports so the team can passively view production traffic. The team will need authorisation to be near and operate on the production ICS. All the testing activities should be closely correlated with the ICS administrator to ensure the process remains functional. The asset owner will ultimately have to understand and accept the risks of testing on the production ICS.

Potential impact of the testing

The potential for disrupting the ICS is great. The team's activities should, therefore, be mostly passive, such as monitoring traffic and events, as opposed to manipulating a data stream. If the team wants to test an attack vector, they should work with the ICS administrator to see whether the target component can be temporarily isolated from the production system. If the answer to this question is no, then this attack vector should be left until the answer becomes yes.

Description of the types of results to expect

The team will be more restricted in what they are able to test on a production ICS. Therefore, the results of this assessment will include an evaluation of the protections that limit an attacker from gaining access to the ICS. For example, the team may evaluate the electronic perimeter of the ICS and provide feedback on how this important boundary can be strengthened. The team also may spend time looking at functional aspects of the ICS such as the patching procedure and replication mechanisms.

Level of security assurance to expect

The team should be able to provide a good analysis of the protections that limit an attacker from reaching the ICS. However, the team will not be able to test much of the ICS because this testing could cause the system to crash. The information provided by the team will be valuable, but the focus will be one layer above the ICS processes and equipment.

Example


The team may have previously evaluated a given ICS in the laboratory and is now performing an onsite assessment to validate the attack scenarios established in the laboratory. In this case, the team will try to establish how an attacker could reach the vulnerabilities discovered in the laboratory given the infrastructure (network layout, firewall rules, etc.) at the location.

Conclusion

Onsite cyber security assessments of an ICS play an important part in the overall security of an installation. As long as the asset owner and assessment team understand the implications of this test, they can focus on what will provide value without impacting the process. This test will definitively answer the question of what an attacker could do at this site. This type of assessment could be employed as a follow-up to a laboratory assessment or when no other way is available to test this ICS (maybe there is no backup or spare system).

Decision Criteria

As illustrated in figure 11, a careful assessment of every production site should be performed.



(Figure 11: Production assessment decision graph)

End-to-end penetration assessment

An end-to-end penetration assessment is one in which the goal of the effort is to gain an understanding of how far an attacker could reach. However, the information required to complete this analysis can be obtained piecewise. For example, if the assessment team demonstrates an exploit that allows them to compromise a DMZ server (victim X) from an attacker box (attacker A) on the corporate LAN and they have shown an exploit for another DMZ server (victim Y), it is reasonable to infer (provided firewall rules are not preventing such) that attacker A could gain remote control of victim Y. Creating the single exploit that chains these events together is busy work for the assessment team that in the end does not provide any more useful information to the organisation funding the assessment. The simple fact that the two exploitable vulnerabilities exist is enough to demonstrate the possibility. There are never enough testing hours; the assessment team should maximise what can be done within the assessment budget.

Also, it may be less effective for the ICS cyber security team to perform internet-in penetration tests (assuming that the control network is protected behind several security zones from the internet). It is typical for the test plan to assume that the attacker has already gained a network presence on a corporate host (search the internet for the quantity of Web browser, e-mail and common service vulnerabilities to understand why). Therefore, the ICS cyber assessment team can focus on those items that are unique to the ICS domain. Many companies in industry specialise in performing internet-in penetration tests of corporate networks. These companies are more suited to perform this work. Therefore, it may be less effective for the ICS cyber security team to attempt any social engineering, client-side and Web application attacks.

Instead, they should focus on the interconnectivity of the electronic zones related to the ICS and the processes and protocols implemented on the ICS.

Pros: Performing an end-to-end assessment may provide an asset owner with a level of confidence on how vulnerable the installation is to a cyber attack. Any gaps where the cyber team is unable to make a network transition from one network zone to the next are used as good examples to secure other network boundaries.

Cons: It may be possible that many testing hours are spent working on areas that are not directly related to the ICS. Corporate network security is not a new topic. There are countless internet references as well as companies with the information and tools to harden this electronic zone. Developing and chaining exploits that allow penetration all the way from the internet to the ICS waste valuable time that could be spent examining and hardening the ICS.

Resources and prerequisites for testing

To perform an end-to-end test requires that the team be given permission to attack all parts of an organisation, ranging from the corporate hosts, DMZ servers and the ICS. Also, because the team would be attempting an attack on a production ICS, there are special considerations for this task.

Potential impact of the testing

The impact on the ICS is the same for a production system. The impact on areas such as the corporate network is less because there is less risk that a crash in one or more hosts will have the effects that a crash in the ICS may have. However, it is likely that the team will cause processes and potentially whole servers to become unresponsive. Damage to databases is also possible as these functions may be used to bridge security zones.

Description of the types of results to expect

The results of this assessment will provide a good indication of locations in the infrastructure (corporate, DMZs and control network) where vulnerabilities exist to exploit. The assessment team will be looking for quick wins that will help them transition from one network zone to the next.

Level of security assurance to expect

An end-to-end assessment may do very well to identify vulnerabilities in the corporate and DMZ networks, but this assessment will likely spend too much time on the first two zones to be effective at testing the ICS. Also, those companies who are experienced penetration testers may feel at home in the corporate and DMZ networks but have little experience on an ICS. Therefore, they may not perform much work in this area or may initiate tests that cause unforeseen crashes in the ICS.

Example

A utility may want to perform an end-to-end assessment to find weak spots requiring improvement. This utility may be more concerned with whether an attacker could reach the ICS than understanding what attackers could do once they have access to the control network.

Conclusion

An end-to-end assessment is effective when the goal is to understand if an attacker could reach the control network.


This tests the effectiveness of the ICS perimeter defences at preventing access to vulnerabilities in the ICS.

An end-to-end assessment is effective when the goal is to understand if an attacker could reach the control network. However, if the goal is to test the ICS, the assessment team should assume that the attacker will find a way to reach the control network.

A lab assessment focuses on the ICS software; a production system assessment focuses on the ICS network and host security without putting it at risk. A penetration assessment tests the ability to reach the ICS. Together, these three methods provide complete coverage of the ICS's security.

Decision Criteria

Factors that influence the value of a penetration test are the reliance on perimeter protection and the corporate IT's security posture. A penetration test should be included if the company or ICS vendor's position is that vulnerabilities in the ICS are irrelevant because the perimeter will protect them. A penetration test will provide less value if the corporate IT group has security assessments performed on its networks. This logic is illustrated in figure 12 below.



(Figure 12: Penetration assessment decision graph)

Component testing

Component testing is testing pieces of an ICS separately from the rest of the system. These tests usually work with the target component isolated (disconnected) from the rest of the ICS.

An example of a component test is a PLC, RTU, HMI application, or database that plays a significant role in the ICS.

Pros: Component tests allow the assessment team to work on a piece of the ICS without the liability of causing cascading problems in the rest of the system. This type of test allows the assessment team to explore how this component operates and whether it has inherent security problems.

Cons: Isolating a component from the rest of the ICS means that the team will not be able to observe how the target interacts with the rest of the system. This situation may not exercise code segments that could be vulnerable to a cyber attack. The team could create their own client to interact with the target component, but this effort takes considerable time as the protocol may be extensive (e.g., ICCP). Also, the system vendor may not publish the documentation required to directly communicate with the component because the system vendor intended it to be used in conjunction with the rest of the ICS.

Resources and prerequisites for testing

The team will need access to the component they intend to test. This component will not be available for service for the course of the testing and potentially afterward if it is damaged during testing. The team will benefit by understanding how this component is used in the larger ICS. For example, if the team can capture network traffic to and from the component before it is isolated from the rest of the system, they are able to search for normal traffic sequences that can be tested offline. An RTU may be capable of talking over many different protocols; therefore, the team should focus on the one that is being used at the target installation.

Potential impact of the testing

The target component will be offline for the duration of the test. Therefore, crashes in this component will have no effect on the rest of the ICS. However, if this component is part of a redundant pair, a dependency problem may exist because the target component may be left in a state where it cannot be quickly placed back in service.

Description of the types of results to expect

Component testing may uncover vulnerabilities that can be accessed in a direct attack. If the configuration of this component does not allow it to authenticate a user, this would be an example of a problem that is inherent in the component. Vulnerabilities that stem from interactions with other components may or may not be exposed in this type of test due to the target being isolated.

Level of security assurance to expect

Component testing will be most effective in determining whether this component is vulnerable to a direct attack. In this scenario, the attacker has gained network access to this component and is trying to attack it using the network footprint of the component.

Example

A good example of a component test is when the assessment team attacks a controller, RTU, or PLC that is not connected to any real input/output (I/O) lines and is also isolated on the IP network.

Conclusion

Component testing can be a valuable task if enough information is available. The biggest hindrance for this type of test is that the assessment team will not see how other components communicate with the target. Therefore, much of the component's functions (potential attack vectors) will be dormant during the test.

Component testing should be used to first eliminate vulnerabilities inherent in the component. Then vulnerabilities that stem from interactions with other components can be identified in a lab assessment.

Decision criteria

Components should be thoroughly tested before they are released. ICS owners can assess the security of individual components by performing component tests. This is cost productive if a particular component:


- has high exposure to attack;
- is critical to ICS operations or safety, or
- has a poor performance (reliability) record;

and

- the component's vendor fixes the identified bugs and vulnerabilities, or
- the owner is able to mitigate them.

A component test may also be the most cost efficient variant of a laboratory assessment if a full backup or development system is not available, but individual components are. There still must be a high level of assurance that the component vendors will fix the reported vulnerabilities. If this is the case for only a subset of components, a component test may be a better assessment target than a laboratory test. If a complete assessment system is available, the individual components may be targeted to obtain more representative results than if the component was disconnected.

Figure 13 shows the simplest decision criteria for performing a component test.



(Figure 13: Component test decision graph)

Technical documentation review

A technical documentation review examines an ICS by looking over documents such as system inventory, architecture diagrams, process diagrams, procedures and process documents.

Pros: One of the benefits of conducting this type of system review is that this task does not affect the production equipment, which means this assessment can be safely conducted on a production system. Documents, such as the network architecture diagrams, help the team identify the electronic perimeter of the ICS. Other documentation, such as the procedure and process data, help the team identify areas where the process could be improved. A document review can help the team identify attack vectors for an actual test and also see if the documentation includes sensitive information that should not be available to the public.

Often, the documentation for an installation is either out-of-date or does not exist. The documentation review can help identify out of date documentation such as inaccurate network diagrams. Insufficient secure configuration documentation can also be discovered during a documentation review. This is a problem with ICS, because the industry trend is to configure a system once and try to minimise future changes to the initial configuration. Because ICS are known for their fragility, ICS administrators follow the old adage, 'If it isn't broke, don't fix it'. Unfortunately, the reality is often that the initial configuration is dependent on the integrator's knowledge and attention to detail because there is no configuration procedure or documentation.

Cons: While a documentation review provides the team with background information for the ICS, it does little toward identifying/fixing cyber security vulnerabilities.

As far as testing for cyber security vulnerabilities, the types of documents that a vendor or an organisation creates will do little toward identifying vulnerabilities because the goal of this documentation is to educate the reader on how to use the system correctly. Cyber security vulnerabilities are often found by using the system in an abnormal manner that the designer did not intend. Vendor documentation rarely goes into a level of detail that is useful to the attack team (the format of data on the wire) because many ICS protocols are proprietary.

Resources and prerequisites for testing

The team will need the current system documents to perform these reviews. Empirical testing results have revealed that these types of documents are usually immature, if they exist. If the goal of the document review is to improve on the ICS process, the review team should consist of ICS engineers rather than cyber security personnel, because there will be less cyber security specific material in this data.

Potential impact of the testing

Reviewing documents will not have an impact on the ICS, which means this effort is safely achieved onsite at a production system installation. It also could be performed before or after a cyber security assessment.

Description of the types of results to expect

A technical documentation review identifies areas where the architecture or process may be improved (assuming that the system is currently functioning).

Level of security assurance to expect

A document review will provide little to no security assurance. There is value in performing this task, but the effort is not much different than performing a cyber security assessment.

Example

The assessment team may perform a document review prior to performing an actual cyber security assessment. This effort helps the team identify that the target system uses X technology, Y protocols, or Z authentication. Having this type of prior information may help the team be prepared for an actual assessment.


Conclusion

A technical document review can be an effective tool if the goals for the task are to prepare for a cyber security assessment or to improve the process. However, this effort will not be able to identify vulnerabilities in the hardware and software that make up the ICS. Instead of a formal task, an assessment team may perform 'need-based' document reviews as they go about a test. For example, they may review architecture diagrams prior to beginning an assessment and reference other documentation when they encounter pieces of the ICS they want to explore further.

Decision criteria

The assessment team must conduct a production assessment without putting the production system at risk. This means that only non-intrusive methods can be used. Production assessments should include documentation, functionality and configuration reviews, along with

staff interviews. These methods must be used to gather information that cannot be obtained by scanning or attacking the system. Documentation decision criteria are shown in figure 14 below.



(Figure 14: Technical documentation decision graph)

Functionality and configuration review

The assessment team should examine the ICS by validating the functionality and checking the configuration of the system. The team incrementally checks the functionality of devices and features of the ICS. Also, the team checks the configurations of many of the ICS components.

Pros: This effort will help the assessment team understand the ICS's unique requirements and characteristics. This activity could identify areas where the process can be optimised. This is the only way to assess and secure the production system components and network.

Cons: The system is not actually tested, but intrusive testing is not an option on production ICSs.

Resources and prerequisites for testing

In order to validate features of the ICS, the team will need access to the target system. For this task to be effective, the team either will have to include an ICS engineer or possess ICS experience prior to this effort. This requirement is sort of a reverse order problem as a cyber security assessment of an ICS should answer security posture questions rather than educate the assessment team about the target system.

Potential impact of the testing

Validating the functionality of an ICS requires the team have access to the system and permission to execute different functions. Unless there are inherent problems in the system, this effort should have little impact on the operations of the ICS other than the distraction of running abnormal functions.

Description of the types of results to expect

This effort may identify areas where the process can be optimised. Examining the device configurations may also identify extraneous services or features that can be safely disabled.

Level of security assurance to expect

High. This process can be used to evaluate the configuration of ICS hosts, network equipment and ICS equipment. Any other non-production assessment of host and network configurations will probably not be representative.

Example

The team could validate the functions of the ICS FEP and check the configuration files.


Conclusion

The assessment team should examine the ICS by validating the functionality and checking the configuration of the system. The team incrementally checks the functionality of devices and features of the ICS. Also, the team checks the configurations of many of the ICS components.

This effort will help the assessment team understand the ICS's unique requirements and characteristics. This activity could identify areas where the process can be optimised. This is the only way to assess and secure the production system components and network.

Decision criteria

The assessment team must conduct a production assessment without putting the production system at risk. This means that only non-intrusive methods can be used. Production assessments should include documentation, functionality and configuration reviews, along with staff interviews. These methods must be used to gather information that cannot be obtained by scanning or attacking the system.



(Figure 15: Functionality and configuration review decision graph)

Staff interviews

The team could formally interview the ICS staff (e.g., process engineers, operators, vendors, integrators, developers, owners and even managers). The goal of these interviews would be to gain further understanding and insight into the processes and procedures of the ICS.

Pros: This effort will help the assessment team understand more of the ICS centric aspects of the ICS and identify areas where the process can be optimised. The team may also learn empirical details about the ICS that would not be found in the system documentation.

Cons: Staff interviews may not have a high level of assurance.

Resources and prerequisites for testing

The team will need access to the staff so they can ask questions and some prior knowledge of ICSs in order to ask the correct questions. This process will greatly increase the assessment team's knowledge of ICS issues and the ICS team's security understanding.

Potential impact of the testing

This task will not have an impact on the operations of the system.

Description of the types of results to expect

These interviews may provide limited insight into the format of data on the wire, though they may capture security concerns staff have already identified but that management have not addressed. It is common for ICS documentation and marketing to present a different view of the system than the way things really are implemented. Even system integrators have limited insight when it comes to the implementation details of the system because they do not work at this level. Rather, the system integrators understand the processes to get the system up and working properly.

Level of security assurance to expect


Interviewing key ICS staff will educate the team on the ICS and the staff's security knowledge, but these interviews will not have high assurance unless combined with access to system components.

Example

The team could interview the system operators to try to uncover any security weaknesses in the way the system is designed or operated.

Conclusion

Interviewing key ICS staff should be part of a production assessment and the documentation and configuration review processes.



(Figure 16: Staff interviews decision graph)

Risk assessment

Risk analysis is used 'to determine whether an asset is protected and to what level. Risk assessment is the quantitative or qualitative process of performing this analysis'^r. In general terms, a cyber security risk assessment is a mathematical way to estimate the likelihood that a system can be attacked using cyber means. Risk assessments often are associated with metrics, models and graphs. The idea is that an analyst identifies the threats to an ICS from observations and by checking configurations and then contrasts these threats against the controls that are in place to protect the system. Each of the attack scenarios is assigned a probability rating so that an end value may summarise the risk to the ICS. Several organisations have created guides, available on the Internet, for assessing the risk to an ICS.^s

r. Matt Bishop, *Computer Security: Art and Science*, Addison-Wesley, 2003.

s. www.controlglobal.com/articles/2005/191.html
ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=1270402
www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6V3P-4P59S0G-1&_user=10&_rdoc=1&_fmt=&_orig=search&_sort=d&view=c&_acct=C000050221&_version=1&_urlVersion=0&_userid=10&md5=0a03963baa4865f3023b25901185f016

Pros: A risk assessment prompts people to think about the ICS from a security perspective. This effort uncovers problems in configurations such as firewall rules or unneeded services. Also, this effort requires less time than an actual penetration test.

Cons: The metrics and conclusions that result from this effort will be educated guesses rather than empirical testing results.

Resources and prerequisites for testing

To perform a risk assessment, the team needs ICS information that may include firewall rule sets, host patch levels and network diagrams. The team may need limited time with the ICS administrator and operators. The team either will have to adopt one of the published risk assessment methodologies for an ICS or construct one of their own.

Potential impact of the testing

A risk assessment will not have an impact on the actual system. All of the team's efforts will be passive.

Description of the types of results to expect

The product of a risk assessment will be a set of metrics, graphs and values that attempt to summarise the risk to the ICS from a cyber attack.

Level of security assurance to expect

Because a risk assessment is based on statistics, the results may or may not represent the resilience of the system to cyber attack. Goodhart's Law states, 'any observed statistical regularity will tend to collapse once pressure is placed upon it for control purpose.'^t Similarly, the results of a risk analysis may be invalid because the results will not account for the irregular way in which vulnerabilities are uncovered and exploited by an attacker.

Example

The assessment team could perform a risk assessment of an ICS (production or offline) based on the architecture diagrams, configuration files and comments from the system integrator.

Conclusion

A risk assessment numerically determines the likelihood that an ICS can be attacked. This method requires less time and resources than a penetration test. Also, this method can be safely accomplished on a production system because the team's efforts would be passive in regard to system operations. The results of a risk assessment may or may not be a good indicator of the security of the system. For example, a risk assessment may determine that many cyber attacks are not viable due to the presence of a firewall or network security appliance (Intrusion Detection System/Intrusion Prevention System). This determination could be a placebo if the protection device is mis-configured, or does not have adequate signatures. Therefore, risk assessments are appropriate for certain situations but should never replace manual testing if the goal is to assess a system for vulnerabilities.

A risk assessment is often part of a vulnerability assessment.

Some form of risk assessment must be performed in order to create attack targets for the assessment plan. This process involves identifying possible attack paths into the ICS (likelihood) and high consequence attack targets (impact).

t. K. Alec Chrystal and Paul D. Mizen, *Goodhart's Law: Its Origins, Meaning and Implications for Monetary Policy*, 2001.

Vulnerabilities can be ranked and prioritised based on their likelihood of being exploited and the associated potential consequences. The easiest way to rank identified vulnerabilities is by their CVSS Version 2 scores.

Assessment methodologies conclusion

When performing an ICS assessment, one of the first decisions that must be made is the target of the assessment. The choice is any combination of the following three categories of targets:


- ICS products (lab assessment);
- ICS network (production assessment);
- ICS perimeter (penetration test).

All three assessment methods are required to perform a thorough assessment of the ICS's security risks. Budget is the primary limiting factor. An ICS owner should first assess his ability to remediate the source of identified vulnerabilities. A vulnerability assessment is of little value if the identified vulnerabilities are not remediated.

In general, ICS owners have the most control over their own networks. A network, or production, assessment should be top priority in this case. A penetration test may be the next priority for the same reason.


A lab assessment is important for assessing the security posture of the ICS products. The value is dependent on whether the vendor fixes the vulnerabilities identified in the ICS products, or the owner's ability to mitigate or detect attacks against them.

The assessment can then be further refined, if desired. A lab assessment may only focus on a subset of the ICS products. When ICS products are evaluated individually, this is called component testing. The most significant difference is that there is no ability to determine the effects of exploits on the behaviour of the system as a whole. As a good practice, components should first be tested individually and then as part of an integrated system. The complete lab assessment decision graph is displayed in figure 17 below.



(Figure 17: Lab system assessment decision graph)

The assessment team must conduct a production assessment without putting the production system at risk. This means that only non-intrusive methods can be used. Production assessments should include documentation, functionality and configuration reviews, along with staff interviews. These methods must be used to gather information that cannot be obtained by scanning or attacking the system. This decision graph is shown in figure 18 below.



(Figure 18: Production system assessment decision graph)

Budget is generally the determining factor on the size of the assessment scope. If a facility is given a large enough assessment budget, the assessment process can be conducted in the following general order, or in parallel:

1. Operational risk assessment
2. Lab assessment
3. Component testing
4. Technical documentation review
5. Functionality and configuration review
6. Production assessment
7. Technical documentation review
8. Staff interviews
9. Functionality and configuration review
10. End-to-end penetration assessment

With a less than optimal budget, any subset of the above assessment methods may be utilised. INL ICS assessment experience has found the methods described below to effectively identify vulnerabilities in ICS installations.

INL production ICS assessments (i.e. on-site assessments) concentrate on the aspects of the ICS that the system owner is able to control, such as secure configurations and layers of

defence. The assessment team only performs penetration testing on disconnected backup or development systems.

The ICS network administrators review and discuss production network diagrams, ACLs, firewall rules and IDS signatures with the assessment team. They can then perform hands on assessments of ICS and network component configurations together. This includes a review and tour of the production system to help identify through documentation, observation and conversation any possible security problems with the production system and network configuration without putting the operational (production) system at risk. This is a learning opportunity for both the assessment team and the asset owner personnel.

The INL has had a lot of success in assessing ICS security while educating vendors and owners on how they can make their systems more secure.

The value of a security assessment can be measured by the increased security (or decreased risk) that results from it. The assessment methodologies can be ranked by the relative decrease in risk that results from mitigating identified vulnerabilities. This is a function of the expected number of vulnerabilities uncovered and the potential impact from exploiting them. Only mitigated vulnerabilities are relevant.

The level of assurance between the different methodologies that the assessment team was able to identify the most likely to be attacked vulnerabilities with the highest consequence is potentially more dependent on the quality of the assessment team than the methods used. As a general rule, however, white-box testing provides the highest assurance that the greatest number and highest consequence vulnerabilities will be identified.

The cost of exploiting known coding and configuration weaknesses is most likely not worth the added assurance, especially when the potential vulnerability can be remediated at a lower cost. In addition, failure to exploit a known weakness may be more of a reflection on the assessment team and current attack techniques.

The most value from the security assessment can be attained by meeting the following objectives:

- Prioritise assessment targets by potential consequences and exposure to attack based on knowledge of the system.
- Employ methodologies that have a high assurance of identifying vulnerabilities that can be mitigated.
- Production network and host assessments are generally most beneficial because the asset owner has the ability to mitigate these vulnerabilities.
- After network and host weaknesses have been secured to the level allowed by the ICS, component, lab, or end-to-end assessments will be more valuable.
- If an ICS vendor does not provide patches for identified vulnerabilities, component and laboratory assessments of their products provided limited value. These methods can give an indication of the system's security level and identify components that require stronger protection and monitoring.
- Promote identification of the most security weaknesses possible, and staff education, by making documentation and staff available to the assessment team.
- Maximise the number of vulnerabilities effectively mitigated by employing the assessment team's assistance in implementing and validating mitigations during the assessment.

- Employ non-intrusive methods such as documentation, functionality, configuration reviews and staff interviews for production system assessments.
- Make existing backup components and systems available to the assessment team for activities that cannot be performed on the production system, such as information gathering, experimentation, validation, etc. This increases assurance at no extra cost.
- Look for opportunities to encourage vendors and other users to assess and secure ICS products. Evaluate product security and vendor's willingness to respond to vulnerability notifications before purchasing new systems.

All assessment methodologies can be balanced to provide the most value from ICS assessments by utilising existing resources to target ICS components that are likely to have the most vulnerabilities that can be mitigated. The quality of the assessment team is most important, but as long as the ICS staff is able to relay ICS requirements, ICS knowledge may not be required. The assessment team must be able to adjust their methods and recommendations to account for non-typical requirements and roadblocks in the ICS environment. ICS staff can work with security experts to identify vulnerabilities without putting the system at risk, rank the risk they pose to the system, and identify mitigations and defence in depth solutions that can be implemented without impacting required functionality. The most effective ICS assessments are coordinated efforts between security experts and ICS administrators who understand their systems well.

Conclusion

ICS owners and operators may employ a cyber security assessment to find out whether their system is vulnerable to a cyber attack. This effort can take many forms including a laboratory, production, or modelling assessment. It is imperative that ICS owners and assessment teams understand the potential impact to the ICS from the testing operations. Any activities that may put the production system at risk should be performed on an offline system so that failures in the ICS will not impact the business or safety of an installation.

This guide covered many aspects of a cyber assessment including the planning, execution and reporting phases to assist those responsible for procuring or facilitating a cyber assessment of an ICS. Important planning, execution and reporting activities are highlighted below.

Planning:

- Ranking ICS components and functionality by potential consequences due to loss of required functionality, data integrity or access control (worst-case consequence analysis).
- Considering the goals and focus of the testing and the amount of information provided to the assessment team. (Remember in most cases it is preferable to assume a worst-case scenario and to provide the testers with as much information as they require, assuming that any determined attacker would already have acquired this. Likewise, ICS owners should be wary of wasting effort on Internet to Corporate Networks tests.)
- Securing the ICS applications, hosts and networks as much as possible, noting security holes that cannot be fully mitigated due to ICS operational requirements. (Eliminating the 'low-hanging fruit' and identifying the most important security goals and obstacles can foster a more valuable security assessment because the assessment team will be forced to search deeper for vulnerabilities. In addition, ICS system and network administrators will also be better prepared to discuss the most important security goals and obstacles, unique to their ICS installation and operational requirements, when creating the test plan.)
- Choosing and appointing the cyber security assessment team, including determining the rules of engagement and execution of a non-disclosure agreement.
- Creating the test plan, remembering not to set too detailed parameters and to allow the assessment team to use their initiative so as to maximise the number of vulnerabilities that can be discovered.
- Determining, in conjunction with the assessment team, the most appropriate assessment methodology for the security test.
- Obtaining vendor involvement, so as to remediate ICS product vulnerabilities if necessary.

Execution:

- Execution of the ICS cyber assessment should be a collaborative effort between the ICS and the assessment teams.
- Prioritisation of ICS system targets can guide the assessment team toward more significant vulnerabilities. Operational requirements and security obstacles can be taken into account by the assessment team when preparing mitigation recommendations.

- Interaction between the ICS and assessment teams during the assessment not only helps guide the assessment team, but facilitates knowledge transfer to the ICS administrators as well.
- Return on investment can be maximised by utilising the assessment team as consultants in securing the ICS, not just breaking it.

Reporting:

- The asset owner should define and require the desired level of reporting during the planning stages of the assessment.
- Reporting requirements should be solution-oriented and tailored to the unique ICS.
- Documentation of potential consequences due to successful exploitation of vulnerabilities should be specific to the environment, which may include the vulnerable ICS application, hosts and networks. ICS administrators can then determine the criticality of unauthorised access or DoS to the affected component.
- During the assessment, ICS administrators and managers should discuss each finding with the assessment team and utilise their expertise to implement as many mitigation techniques as possible. This can reduce the need for validation testing of mitigations for assessment findings.
- Recommendations should address system requirements and ICS administrators responsible for implementing them should have the opportunity to make sure they understand how.
- The assessment team can also help the ICS owner work with the ICS vendor to remediate vulnerabilities in the ICS products and define ICS product security requirements for future procurements.

Maximising return on investment

A key goal of this guide is to help asset owners maximise the return on investment when undertaking systems testing for security vulnerabilities. Return on investment can crudely be measured by the increase in security (decrease in risk) divided by the cost of the additional security.

In a simplistic view, the best assessment methodology is the one that promises the highest vulnerability reduction at lowest cost. The benefit from a vulnerability assessment is proportional to the number of vulnerabilities that are identified for remediation and is therefore dependent on the asset owner's ability to mitigate the identified vulnerabilities.

The asset owner should ensure that the assessment team provides adequate vulnerability details and mitigation information for the ICS administrators or vendors to efficiently and effectively remediate each security weakness. This will minimise duplication of effort in vulnerability identification.

A collaborative assessment is most valuable, because ICS staff gain security knowledge directly applicable to their system and can take advantage of the assessment team's knowledge to immediately mitigate many vulnerabilities as they are identified.

The most cost beneficial reporting method is to clearly and concisely document vulnerability details needed to assess potential impact, vulnerability location and remediate information.

Vulnerabilities should be prioritised by reduced risk per remediation cost. Risk can be measured using the CVSS v2 vulnerability scoring method.

Glossary

Acronyms

ANSI	American National Standards Institute
ASCII	American Standard Code for Information Interchange
CIA	Confidentiality, Integrity and Availability
CISSP	Certified Information Systems Security Professional
CVSS	Common Vulnerability Scoring System
DCS	Distributed Control System
DHS	United States Department of Homeland Security
DMZ	Demilitarised Zone
DoS	Denial-of-Service
FEP	Front-end processor
FTP	File Transfer Protocol
GIAC	Global Information Assurance Certification
GSEC	GIAC Security Essentials Certification
HMI	Human-machine Interface
HTTPS	HyperText Transfer Protocol over Secure Socket Layer
I/O	Input/output
ICCP	Inter-Control Centre Communications Protocol
ICS	Industrial Control System
ID	Identification
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
INL	Idaho National Laboratory
IP	Internet Protocol
ISO	International Organisation for Standardisation
IT	Information Technology
LAN	Local Area Network
LM	LAN Manager
NMap	Network Mapper
PCS	Process Control System
PLC	Programmable Logic Controller
PPE	Personal Protective Equipment
RTU	Remote Terminal Unit

SPAN	Switched Port Analyser
SQL	Structured Query Language
TCP	Transmission Control Protocol
VPN	Virtual private network
WAN	Wide area network
XSS	Cross-site scripting

Nomenclature

Black-box testing	Black-box testing takes an external perspective of the test object to derive test cases. These tests are usually functional. The test designer selects valid and invalid inputs and determines the correct output. There is no knowledge of the test object's internal structure.
Clear-text	Clear-text is the form of a message or data that is immediately comprehensible to a human being without additional processing.
Full disclosure	In computer security, full disclosure means to disclose all details of a security problem which are known.
Fuzz testing, fuzzing	Fuzz testing is a software testing technique that provides random data (fuzz) to the inputs of a program.
LanMan hash	LM hash or LAN Manager hash is one of the formats that Microsoft LAN Manager and Microsoft Windows versions previous to Windows Vista uses to store user passwords that are fewer than 15 characters long. This type of hash is the only type of encryption used in Microsoft LAN Manager (hence the name) and versions of Windows up to Windows Me.
Man-in-the-middle	In cryptography, man-in-the-middle is a form of active eavesdropping in which the attacker makes independent connections with the victims and relays messages between them, making them believe that they are talking directly to each other over a private connection when in fact the entire conversation is controlled by the attacker. The attacker must be able to intercept all messages going between the two victims and inject new ones, which is straightforward in many circumstances (for example, the owner of a public wireless access point can in principle conduct man-in-the-middle attacks on the users).
Pentest	Penetration testing is the security-oriented probing of a computer system or network to seek out vulnerabilities that an attacker could exploit. The testing process involves an exploration of all the security features of the system in question, followed by an attempt to breach security and penetrate the system.
Plaintext	Plaintext is the information that the sender wishes to transmit to the receivers. Before the computer era, plaintext simply meant text in the language of the communicating parties. Since computers, the definition has been expanded to include not only the electronic representation of text, such as email and word processor documents, but also the computer representation of speech, music, pictures and videos.
Port mirroring	Port mirroring is used on a network switch to send a copy of all network packets seen on one switch port (or an entire VLAN) to a network monitoring connection on another switch port. This is commonly used for network appliances that require monitoring of network traffic, such as an intrusion detection system.
Risk assessment	Risk assessment is the determination of quantitative or qualitative value of risk related to a concrete situation and a recognised threat (also called hazard).

Security assessment	An explicit study to locate IT security vulnerabilities and risks.
Span port	Switched Port Analyzer (SPAN) port is a more common name for port mirroring on a Cisco Systems switch.
White-box testing	<p>A tester knows the internal program structure and its code. The tester can execute each program statement and function; check error handling, etc. This testing involves source code reviews, walkthroughs, as well as design and execution of tests based on the access to the program code.</p> <p>White-box testing requires deeper knowledge of programming languages and technologies than black-box testing.^u</p>

^u Software testing glossary, Testing Types and Activities, BugHuntress, bughuntress.com/analytics/glossary.html

Figures

Figure 1: Black-box versus white-box testing	7
Figure 2: Benefits of black-box and white-box testing combined in grey-box testing	8
Figure 3: Generic IT security goals versus ICS security goals	10
Figure 4: Assessment process flow chart	13
Figure 5: Example of supervisory control LAN attack targets	16
Figure 6: Potential attack vectors between ICS network security zones	18
Figure 7: Sample cyber security assessment report outline.....	26
Figure 8: CVSS Metric groups.....	29
Figure 9: Sample ICS configuration for a laboratory assessment	38
Figure 10: Lab assessment decision graph	40
Figure 11: Production assessment decision graph	42
Figure 12: Penetration assessment decision graph	44
Figure 13: Component test decision graph.....	47
Figure 14: Technical documentation decision graph	49
Figure 15: Functionality and configuration review decision graph.....	51
Figure 16: Staff interviews decision graph.....	52
Figure 17: Lab system assessment decision graph.....	56
Figure 18: Production system assessment decision graph.....	57

Tables

Table 1: Popular network reconnaissance tools	20
Table 2: Open ports on host (IP) example.....	21
Table 3: Example of assessment boundary list	27
Table 4: CVSS v2 base scoring metrics	30
Table 5: CVSS v2 temporal scoring metrics	31
Table 6: CVSS v2 environmental scoring metrics.....	32
Table 7: Vulnerability severity ratings.....	33