Quantitative Security Measures for Cyber Security Assurance

Dennis Holstein
Opus Publishing
Seal Beach, California
holsteindk@ieee.org

Abstract: ISA (International Society of Automation) and others are struggling mightily to develop quantitative cyber security assurance metrics. All metrics proposed to date are at best qualitative and highly subjective. They are certainly without mathematical rigor; and clearly, they do not conform to the guidelines recommended by Andrew Jaquith in his book "Security Metrics - Replacing Fear, Uncertainty and Doubt" [1]. Papers by Kube, Langer and Singer [2,3] presented in the 2008 S4 Conference address the difficulty in developing quantitative metrics. This paper follows the Jaquith model which requires that good metrics must be capable of being consistently measured (without subjective criteria), the measurement of these metrics must be collectable by automated systems, the metrics must be expressed as cardinal number or percentage using at least one unit of measure, and they must be meaningful to the asset owner.

The approach begins by using the ISA's mapping of NIST 800-53 [4] (ISA99WG04TG2, 2009 (Work in Progress) [5]) to the seven foundational requirements as defined in ISA-99.01.01 [6]. The next step, described in this paper, offers the mathematical formulas for quantifying security assurance for the system and for the components of the system under consideration (SuC). These formulas do not rely on minimizing the risk-based based formula $R = (1-P_e)*C*P_o$. Rather, a consequence-based analysis is used to establish weighting coefficients for contributing security mechanisms that are used in the proposed formulas. Thus, we can build a "score card" to collect measurements of the security metrics that can then be used to assess the implemented security assurance of the system under consideration.

Keywords: ISA99, NIST 800-53, Security Metrics, Foundational Requirements

1 Introduction

To paraphrase Lord Kelvin's famous quote, "You cannot improve what you cannot measure." In the forward to Jaquith's book [1], Daniel E. Greer wrote "Security metrics are the servants of risk management and risk management is about making decisions. Therefore, the only security metrics we are interested in are those that support decision making about risk for the purpose of managing that risk."

There are good metrics and there are bad metrics! Jaquith eloquently differentiates the good from the bad, and he establishes the criteria for good metrics.

- They must be capable of being consistently measured without subjective criteria
- The measurement of these metrics must be collectable by automated systems
- The metrics must be expressed as a cardinal number or percentage using at least one unit of measure
- And they must be meaningful to the asset owner

That's our litmus test! Our problem domain is cyber security assurance of an Industrial Automation and Control System (IACS). Researchers are struggling mightily to develop quantitative cyber security measures for this application – to date with little success. So the first question asked is: what makes this so difficult? We know the classic risk-based formula $R = (1 - P_e) * C * P_o$; where R is the measure of risk, P_e is a measure of system effectiveness[1], C is the measure of consequence that the event will result in harm, and P_o is the probability of occurrence (or opportunity) to successfully execute the cyber attack. The difficulty lies in the fact that we are trying to quantify a cyber-security vulnerability which given the lack of empirical data is a low probability, but, if it occurs, the impact measured by the consequence is high.

Torgerson [7] states that quantifying security measures using the probabilities required by the risk equation is impossible. Singer and Kube [3] considered using an approach similar to that used to quantify Safety Integrity Level (SIL). They concluded that the SIL approach was not practical. The arguments offered by Torgerson, Singer and Kube convinced me that using a statistical methodology was not a good approach.

2 A New Approach Offered by ISA

ISA 99 decided to release their standard of Technical Requirements as a series of documents as listed in the References. The first document in the series is a technical report of security technologies, ISA99WG04, ISA-TR99.03.01: Security Technologies [8].

The second document in the series, ISA99WG04, ISA-99.03.02: Security for Industrial and Automation Control Systems – Technical Requirements: Target Security Levels, mid 2009 – Work in Progress [9], specifies the rules to define the zones and conduits of system under consideration – called IACS. Each zone includes automation subsystems and components with a common security assurance objective. Zones may contain other zones to provide logical grouping of IACS components. Containment zones may take credit for the protection offered by the container zone. For example, consider a configuration where Zone A (zA) contains Zone D (zD) and Zone E (zE), and Zone B (zB) and Zone C (zC) are zones external to Zone A. Zone D and Zone E take credit for the protection offered by Zone A with respect to the secured access and use control requested by Zone B or Zone C. Thus, the target security strength of Zone D with respect to access and use by Zone B and Zone C may be expressed as follows:

[1] P_e includes Probability of Detection, Probability of Assessment, and Probability of Interruption / Neutralization.

$$S^{target}_{zD|B,C} = \min [S^{target}_{zA|B}, S^{target}_{zA|C}] \quad \text{Equation 1}$$

The credit hypothesis of Equation 1 is as follows: the security strength of Zone D with respect to the access and use requests by Zone B and Zone C is the minimum of the security strength provided by Zone A conditioned on Zone B (zA|B) and Zone A conditioned on Zone C (zA|C). The basic idea of Eqn.1 has many applications for perimeter defense, compensating security for device constrained components and legacy components. An example is discussed later in this paper.

Zones are connected by conduits to provide the needed communications within the system under consideration as well as the communications to external entities. The methodology used to define the security assurance objectives is specified by ISA in another ISA 99 standard ISA99WG02, Early 2009.

The first document of Technical Requirements also specifies the rules set the target security level for the system under consideration. These rules apply the criteria specified by NIST [4]. This is a good start, but according to Jaquith the resulting target security assurance levels specified by this technique are not adequate and they do not meet the criteria defined above. ISA intends to rectify this situation in the specification of the second document of the Technical Requirements series which will define the rules to quantify the system security assurance level in ISA99WG04, ISA-99.03.03: Security for Industrial and Automation Control Systems – Technical Requirements: System Security Compliance Metrics, late 2009 – Work in Progress [10]. This paper describes a proposed mathematical formulation to quantify these compliance metrics.

3 The Mathematics of Compliance Metrics – A Proposal

Holstein first published his proposal for the mathematics of compliance metrics as a paper for the HICSS-42 conference [11]. For reader convenience, the methodology and mathematics are repeated here for this paper.

3.1 The Need to Account for Time and Event Dynamics

We are concerned with a cyber-attack, or series of attacks, carried out in a manner that is destructive or disruptive of IACS' operations and processes which are networked together by various communications technologies and mechanisms. Figure 1 is a notional description of the relatively slow degradation of the actual security assurance level[2] of the system under consideration over time. The sharp degradation in security results from an adversary initiated event. Security assessments are performed at discrete times, therefore connecting the dots is not a smooth curve, but the trend-line between the time of the installed security and the time of when the adversary initiates an event can be represented as shown in Figure 1.

[2] The actual security assurance level is sometimes labeled the security health of the system.

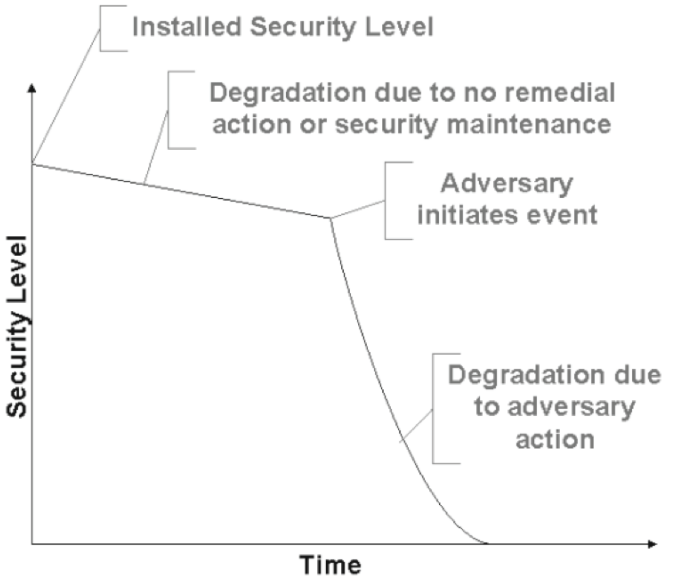

Figure 1 – Notational Degradation of Security Level

Initially, the installed security assurance level of a conduit, zone, subsystem, or system is at or near the design or security assurance level established by the asset owner. There is a natural degradation in the actual security assurance level over time because of a lack of remedial action or security maintenance. For example, if passwords are not changed on a regular basis the effective security level will degrade. If the asset owner is diligent, the security assurance level will plus-up on a regular basis.

At some point in time, which may be selected by the adversary, an event is initiated that causes the actual security assurance level to sharply degrade, resulting in a security breach. The time lapse between the initial event and the security breach may be nearly instantaneous, or delayed, depending on the scenario objectives of the adversary.

3.2 The Math

The security assurance level of the system (S_{system}) may be described as the sum of the weighted security assurance levels of the components ($w_i s_i$).

$$S_{system} = \Sigma_{i=1}^{n} w_i s_i \quad \text{Equation 2}$$

Based on the owner's assessment of risk and evaluation of consequences, w_i (the weighting) and s_i (the security assurance level of the component) are assigned by the asset owner.

It is important to note that s_i is not a statement of probability, and there is no requirement that s_i =[0,1]. It is more like a score. A similar approach is described in the Common Vulnerability Scoring System (CVSS); however, it is not obvious why s_i should be related to a probability of occurrence which is the foundation premise of CVSS. It is more likely that s_i is a measure of the consequence of the failure to adequately protect the system against an adversary induced attack. That is, if s_i is high, the consequence resulting from an attack is low; i.e., a measure of effectiveness.

The target security assurance level of a system can, with considerable effort, be determined for the NIST 800-53 and 800-82 documents. The problem is that there is no insight in the NIST publications as to how an asset owner should allocate system level security assurance target to the component (or subsystem) security assurance levels. Be that as it may, the target value for the system should be greater than the estimated (calculated) value for the system.

$$S^{target}_{system} \geq S_{system} \quad \text{Equation 3}$$

Thus, we have a boundary constraint that can be used.

- The asset owner estimates the target security assurance level of the system (S^{target}_{system}) using NIST 800-53 and 800-82. When required, the asset owner uses "security strength" as defined by Requirement #26 in ISA-99.03.02 to discriminate between security assurance level targets.
- Using the same assumptions and rules set forth in the NIST publications, the asset owner uses Eqn. 2 to calculate the security assurance level of the system by summing the weighted components. This is now the design security assurance level for the system (S^{design}_{system})

If the estimated (calculated) design security assurance level for the system is greater than target security assurance level of the system, something is amiss – probably the assumptions are either not consistent, or, the application of the assumptions has not been correctly performed. In either case the asset owner must redo both estimates.

If the design security assurance level for the system is less than the target security assurance level for the system, the difference is the design margin. Given the uncertainty in all factors of this process, beginning with the uncertainty in the initial risk assessment and consequence analysis as well as the uncertainty in properly applying the assumptions (as discussed above), a design margin on the order of 100% is probably reasonable.

3.3 Weighting Factors are Determined by Quantitative Metrics

Assignment of the weighted security assurance levels of the components ($w_i s_i$) described in Equation 2 requires quantitative metrics. A comprehensive system treatment of these metrics is the challenge for ISA-99.03.03 [10], and allocation to subsystems and components is the challenge for the remaining technical requirements in this series. Because cyber security is highly influenced by the concept of trustworthiness, the metrics for this subject are of particular interest.

Let's take one example: the foundational requirement for Timely Reporting of Events (TRE). Ten of the metrics are:

1. Number of access retries associated with a source address recorded by each access point.
2. Number of rejected access and use certificates requested by each supplicant (requestor) address from each legitimate authenticator.
3. Number of data integrity violations (rejections) related to a source address reported by each industrial automation entity.
4. Number of data confidentiality violations (rejections) related to a source address reported by each industrial automation entity.
5. Number of rejected data packets related to a wireless source address reported by each stateful firewall filter.
6. Number of rejected data packets related to a source address reported by each network resource.
7. Out-of-service time reported by each network resource.
8. Time delay measured from the time of recording an event to the time received by an entity responsible for processing the event.
9. Number of cyber violations associated with trusted human and non-human entities.
10. Time to recover to a state of normal operations from a cyber violation that are associated with trusted human and non-human entities.

Recall the discussion of zones and Equation 1 earlier. The issue was secured access and use control of Zone B and Zone C supplicants (requestors) to components in Zone A – namely Zone D and Zone E components. The objective of the TRE metrics 1 through 10 above is to respond to security violations by notifying the proper authority, reporting needed forensic evidence of the violation, and automatically taking timely corrective action in mission critical or safety critical situations. If adequate means for TRE are implemented, these metrics for an industrial automation system should be available for post-event analysis.

A weighting factor for each component of the SuC may now be assigned based on its functional capability to support a specific metric. If not supported, $w_i = 0$, if supported $w_i = [0,1]$.

3.4 Practical Example – A Chemical Truck Loading Station

Figure 2 shows an example application of a chemical truck loading station (Zone A previously discussed), where a Programmable Logic Controller (PLC) is used to control

the loading of liquid chlorine, a hazardous chemical, from a storage tank to a tank truck. Components of the IACS include:

- Field instruments such as flow sensors and transmitters, valves, motor control for pumps
- PLC system including power supplies, CPU, and interface modules for field signals
- Workstations for operation, supervision, maintenance and engineering
- Communication infrastructure including switches and firewalls

The field instruments and PLC system (Zone D previously discussed) and one operator workstation (Zone E previously discussed) are located in the field, while the workstations for supervision, maintenance and engineering are located in control building some distance (1,000 feet) from the field (Zone B previously discussed). The workstations in the control building are part of the plant network and have other applications running on them (Zone C previously discussed).

Grouping components and systems into zones by the asset owner is based on security requirements that considered both functional requirements as well as physical grouping of assets for the application. The asset owner also considered the consequences of a security breach. For example, overfilling a tank truck with water is of lesser consequence than overfilling a tank truck containing liquid chlorine. A liquid chlorine spill due to overfilling can result in fatalities. Lastly, grouping the components into specific zones reflects the fact that legacy or device constrained IACS components do not have any inherent security capabilities.

For this example, the conduit includes active communication assets that impact the security of the zones it connects. For example, Figure 2 shows a conduit between the control systems security zone and the plant network security zone which contains routers and firewalls. This conduit contains many logical channels: a channel exists between the engineering workstation in the plant network security zone and the PLC CPU in the control systems security zone, channels between the PLC CPU and other workstations in the plant network security zone.

For this example, Table 1 shows results of the asset owners assessment and assignment of weighting (w_i) for each metric (M_i). The rationale for the assignment of weights is:

- Except for the layer 3 switch in the plant network security zone all components deployed provide information sought by every TRE metric except metric #5. These components have the necessary memory and processing power needed to log the forensic data and include a dedicated communication channel to an external security authority to report these data in a timely manner.
- Only the router/firewall has the needed stateful filters and memory/CPU needed to log the number of rejected data packets and report these data to an external security authority.
- Except for TRE metric #5 the operator workstation in the truck loading station has the necessary memory and process power needed to log the forensic data

and include a dedicated communication channel to an external security authority to report these data in a timely manner.

- The field switch, field instruments and PLC provide no security functionality – they rely on compensating security mechanisms in the other components.

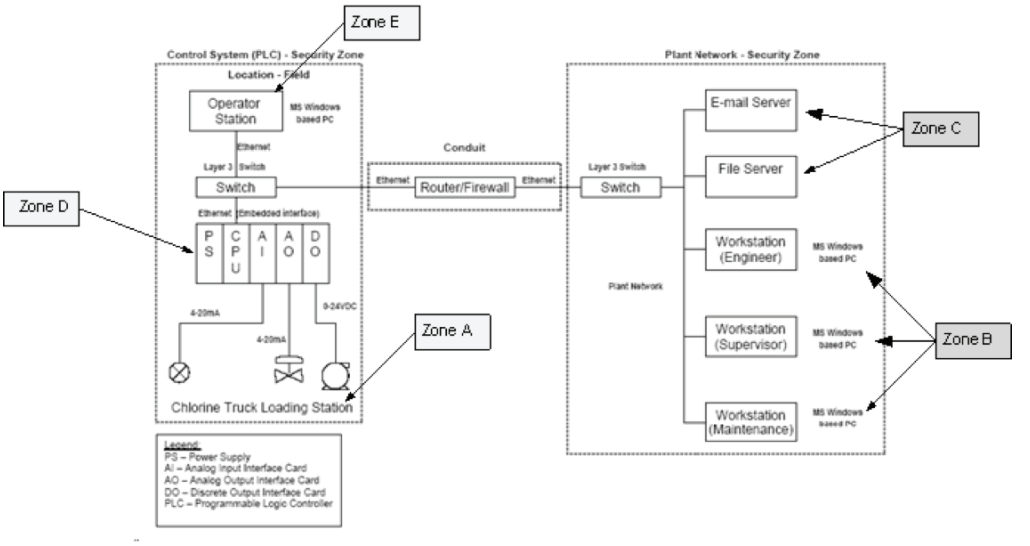

Figure 2 – Chemical Truck Loading Station Example

	M_1	M_2	M_3	M_4	M_5	M_6	M_7	M_8	M_9	M_{10}
Maintenance Workstation	1.0	1.0	1.0	1.0	0	1.0	1.0	1.0	1.0	1.0
Supervisor Workstation	1.0	1.0	1.0	1.0	0	1.0	1.0	1.0	1.0	1.0
Engineer Workstation	1.0	1.0	1.0	1.0	0	1.0	1.0	1.0	1.0	1.0
File Server	1.0	1.0	1.0	1.0	0	1.0	1.0	1.0	1.0	1.0
Email Server	1.0	1.0	1.0	1.0	0	1.0	1.0	1.0	1.0	1.0
Plant Network Switch	0	0	0	0	0	0	0	0	0	0
Conduit (Router/Firewall	0	0	0	0	1.0	1.0	1.0	0	0	0
Operator Station	1.0	1.0	1.0	1.0	0	1.0	1.0	1.0	1.0	1.0
Field Switch	0	0	0	0	0	0	0	0	0	0
Field Instruments	0	0	0	0	0	0	0	0	0	0
PLC	0	0	0	0	0	0	0	0	0	0

Table 1 – Assignment of Security Weights

3.5 Target System Security Assurance Level for the Example SuC

The example includes components selected for a chemical truck loading station and support functions on the plant network. Certified security assurance levels for each component will receive a grade of 1 to 4 based on the TRE requirements specified in ISA-99.03.02 and allocated to the each class of components to be specified in ISA-99-03.xx. ISA is currently struggling with how to grade and certify each class of components. The framework for this approach is accepted; the details are being debated.

The target TRE security assurance level for the system, or , is the sum of all component contributions listed in Table 1. One should note that if a security weight entry in Table 1 is "0" the security assurance contribution of the component for the metric under consideration will be "0" regardless of its certified grade. If the security weight for that component is "0" for all metrics and all foundational requirements, then security assurance for that component must be achieved by compensating mechanisms. This is usually true for legacy components with limited memory or processing power such as the field switch, field instruments and PLC in the example shown in Figure 2.

3.6 Normalization to Set the Target Security Assurance Level

ISA-99.03.02 requires that the sum of the components be normalized to calculate a the target system assurance level. Table 1 includes 11 components and 10 metrics, a combination of 110 contributing members. The non-zero contributing members is a countable set of 57 members which is expressed as $C_{non\text{-}zero} \in C_{members}$ • $C_{non\text{-}zero}$ is then used to normalize the sum of the weighted components. For the example shown in Table 1, if all non-zero members are certified to an ISA-99.03.02 security assurance level of 3, then the normalized target system assurance level would be 3 because all the weights are set to 1. If some of the components in the design are certified to a level other than 3, then the weighted sum is the design TRE security assurance level for the system. Thus, we have the mathematics for a normalized security assurance level of the system that can be used to describe either the target level or the design level.

When deployed, the quantitative metrics, M_i, listed in Table 1 are measured periodically in accordance with company policy. These measurements are then used to calculate the achieved state of the TRE security assurance level of the system. It was this measured state of the system security assurance level that was used for Figure 1.

4 Conclusion and Status

The state of work-in-progress by ISA 99 and related ISA organizations is encouraging. To the best of this author's knowledge, only ISA has focused their attention on developing quantitative security measures for cyber security assurance with due diligence given to the functionality and unique operating requirements for industrial automation systems.

The challenge is to develop and vet a methodology to allocate system security assurance level requirements to subsystems and components. Furthermore, ISA plans to put in place a certification authority with the capability to independently verify that products and services offered by solution providers and manufacturers of industrial automation

systems and components comply with the security requirements specified in the ISA 99 standard.

This is a daunting task that needs contributions from the total community, domestic and international. To this end, ISA and IEC will jointly review and ballot the ISA 99 standards. The market will then judge our success.

About the Author – Dennis Holstein is the Publisher at OPUS Publishing. Under contract to the Electrical Power Research Institute (EPRI), Dennis directed the developed an interoperability test and evaluation laboratory to rapidly prototype design solutions for electrical power delivery substation automation that integrates protection, control and data acquisition on the substation communication network. Dennis is very active in control system security standards efforts including IEEE, ISA, IEC and CIGRE committees and working groups.

References

[1] Jaquith, A. (2007). Security Metrics. In Security Metrics - Replacing Fear, Uncertainty, and Doubt (pp. 22-25). Boston, MA, USA: Pearson Education, Inc.

[2] R. Langner and B. Singer, "SCADA Threat Modeling Using Attack Scenarios", 2008 S4 Proceedings, Digital Bond Press, 2008.

[3] N. Kube and B. Singer, "Security Assurance Levels: A SIL Approach to Security", 2008 S4 Proceedings, Digital Bond Press, 2008.

[4] "NIST SP 800-53: Recommended Security Controls for Federal Information Systems, 2005.

[5] ISA99WG04TG2, "ISA-TR99-01-02: NIST 800-53 Mapping to ISA Foundational Requirements and Security Assurance Levels. Technical Report", (2009 (Work-in-Progress).

[6] ISA99WG02, "ISA-99.02.01: Security for Industrial and Automation Control Systems - Establishing an Industrial Automation and Control Systems Security Program", International Society of Automation, (Early 2009).

[7] M. Torgerson, "Security Metrics for Communication Systems", 12th ICCRTS.

[8] ISA99WG04, "ISA-99.03.01: Security for Industrial and Automation Control Systems - Technical Requirements: Target Security Levels", International Society of Automation, (mid 2009 - Work in Progress).

[9] ISA99WG04, "ISA-99.03.02: Security for Industrial and Automation Control Systems - Technical Requirements: System Security Compliance Metrics, International Society of Automation", (late 2009 - Work in Progress).

[10] ISA99WG04, "ISA-99.03.03: Security for Industrial and Automation Control Systems - Technical Requirements: System Security Compliance Metrics, International Society of Automation", (late 2009 - Work in Progress).

[11] D. Holstein, "A Systems Dynamics View of Security Assurance Issues - The Curse of Complexity and Avoiding Chaos", Hawaiian International Conference on System Sciences, #142, 2009.