INTERNATIONAL JOURNAL OF CRITICAL INFRASTRUCTURE PROTECTION 9 (2015) 52–80

Available online at www.sciencedirect.com

ScienceDirect

www.elsevier.com/locate/ijcip

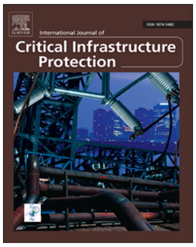

A survey of cyber security management in industrial control systems

William Knowles[a], Daniel Prince[a], David Hutchison[a,*], Jules Ferdinand Pagna Disso[b], Kevin Jones[b]

[a]Security Lancaster, School of Computing and Communications, Lancaster University, Lancaster LA1 4WA, United Kingdom
[b]Airbus Group Innovations, Quadrant House, Celtic Springs, Coedkernew, Newport NP10 8FZ, United Kingdom

ARTICLE INFO

Article history:
Received 17 April 2014
Accepted 4 November 2014
Available online 3 March 2015

Keywords:
Industrial control systems
SCADA systems
Risk assessment
Risk management
Security metrics
Risk metrics

ABSTRACT

Contemporary industrial control systems no longer operate in isolation, but use other networks (e.g., corporate networks and the Internet) to facilitate and improve business processes. The consequence of this development is the increased exposure to cyber threats. This paper surveys the latest methodologies and research for measuring and managing this risk. A dearth of industrial-control-system-specific security metrics has been identified as a barrier to implementing these methodologies. Consequently, an agenda for future research on industrial control system security metrics is outlined. The "functional assurance" concept is also introduced to deal with fail-safe and fail-secure industrial control system operations.

1. Introduction

The number of security-related incidents involving industrial control systems (ICSs) in 2012 was more than five times their 2010 level (197 incidents in 2012 compared with 39 in 2010), according to a report by the Industrial Control Systems Computer Emergency Readiness Team (ICS-CERT) [215]. The rising incident count has been a catalyst for the increased focus on securing industrial control systems.

The default perspective for industrial control system stakeholders has been to view security as a low priority goal, while relying on security through obscurity (i.e., using secrecy in an attempt to ensure security). This technique has seen consistent use, but its success has differed across the three generations of industrial control systems [2]. Security through obscurity largely worked for first generation (monolithic) and second generation (distributed) industrial control systems, which used proprietary and closed-source components and standards, with limited connectivity to non-industrial-control-systems. However, third generation (networked) industrial control systems frequently use open technologies, while connecting to and communicating over other (potentially non-industrial-control-system) networks. This openness has increased the susceptibility to attack, primarily due to greater awareness of industrial control system technologies and their use of standard protocols. Many industrial control systems are often seen as critical infrastructures, making them attractive targets for attack.

The openness of third generation industrial control systems can be illustrated through the use of a reference model (Fig. 1). The lowest level consists of devices that ensure that an industrial control system enters a fail-safe mode when

[*]Corresponding author.
E-mail address: d.hutchison@lancaster.ac.uk (D. Hutchison).

http://dx.doi.org/10.1016/j.ijcip.2015.02.002
1874-5482/© 2015 Published by Elsevier B.V.

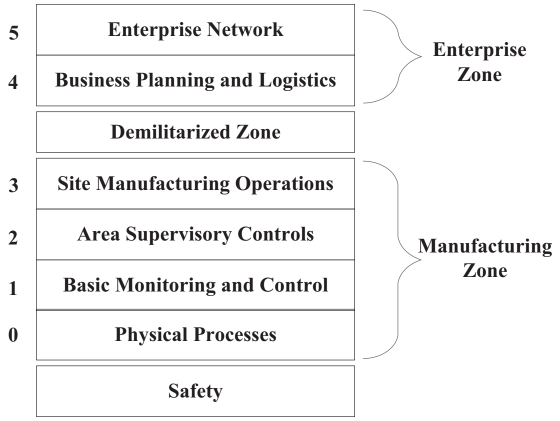
5
Enterprise Network
4
Business Planning and Logistics
Enterprise Zone
Demilitarized Zone
3
Site Manufacturing Operations
2
Area Supervisory Controls
1
Basic Monitoring and Control
0
Physical Processes
Manufacturing Zone
Safety

Fig. 1 – Industrial control system reference model (adapted from [62]).

dangerous conditions occur. Layer 0 includes sensors and actuators that interact with physical processes (without autonomy). Layer 1 devices monitor and control physical process using the sensors and actuators in layer 0; the devices include programmable logic controllers (PLCs) and remote terminal units (RTUs). Layer 2 handles supervisory and operational functions, and often shares data with layers 3–5; the devices include alert systems and human–machine interfaces (HMIs). Layer 3 is the highest level of what would traditionally be defined as the industrial control system network (i.e., manufacturing zone) and provides plantwide functions. Contemporary industrial control systems contain many information technologies at layer 3, and this layer frequently communicates with business applications at layers 4 and 5. The devices include historians (i.e., databases of time-stamped industrial control system events such as process outputs and alarms), and authentication, authorization and accounting (AAA) services. Layer 4 relies on standard information technologies and provides business administration services, such as enterprise applications (e.g., e-mail servers) and non-critical industrial control functions (e.g., inventory management). Layer 5 consists of the majority of centralized information technology services (e.g., business-to-customer services). The reference model highlights the multitude of interconnections between industrial control systems and information technologies both within and between layers. Furthermore, complexity increases due to the diverse communications media used for these interconnections. Although they may be located within a single facility, it is most common for devices within particular layers to be geographically distributed (e.g., a human–machine interface (layer 2) communicating with one or more programmable logic controllers (layer 1) at remote field sites over the Internet).

A challenge arises in the risk management of industrial control systems because standards and methodologies for traditional information technology systems cannot be applied directly. For traditional information technology systems, the order of prioritized security goals on which these approaches are based is typically confidentiality, integrity and then availability (CIA). For industrial control systems, the priority is generally reversed (AIC), with availability as the primary goal [185] (e.g., a utility prioritizing the continuity of service). There are, however, exceptions to the AIC generalization (e.g., when intellectual property is involved in a manufacturing plant). Issues are further compounded when one considers different subsystems with different goals. For example, does an interconnected corporate network exist as part of the industrial control system or as a distinct entity? Functionally, the corporate network is a traditional information technology system that may mandate many non-standard industrial control system requirements (e.g., for information security); however, its interconnection to an industrial control system provides routes for attack and, furthermore, it may also contain systems with control capabilities.

The European Network and Information Security Agency (ENISA) [77] has extended this debate by providing an alternative definition that maintains that industrial control systems are not ruled by CIA, but by safety, reliability and availability (SRA). Safety, in particular, is an important consideration due to its potential to be negatively influenced by security solutions. These complexities highlight the multi-dimensional nature of industrial control system security, and the challenges of measuring its constituent features.

This has led to a variety of new publications (e.g., standards, guidelines and best practices), legislation and other initiatives with the common goal of increasing industrial control system security. However, criticism can be leveled against this body of work due to the lack of guidance on conducting practical security evaluations. A fundamental reason for this criticism is the scarcity of industrial-control-system-security specific metrics. This claim is substantiated later in this paper based on an analysis of the literature (Section 4).

The availability of a comprehensive and robust set of security metrics is essential for organizations to meet various business objectives. These objectives are outlined in a number of publications (e.g., [45,105,126]); however, in summary, there are three broad uses.

The first is to meet demands from external sources. The quintessential example is the obligations imposed by regulations. Although regulations exist for specific use cases of industrial control systems (e.g., in defense), there are no cross-industry regulations. However, this is changing with the implementation of regulations that target critical infrastructures. For example, the European Union (EU) has issued a Directive on Network and Information Security [75], which is expected to be adopted by 2015. Another example of externally enforced usage is meeting contractual demands; this is typically the case for contracts involving government bodies or high-security activities.

The second use case is to evaluate compliance with standards such as ISO/IEC 27001 for information security. Motivations for compliance can be external (e.g., regulations) and internal (e.g., to improve risk posture).

The third use case is to evaluate the risk posture. Although this may be based on both regulatory and compliance motivations, neither is a prerequisite. Examples of this use case include integrating security during the product development cycle (e.g., to minimize software vulnerabilities), supporting strategic decision making (e.g., enterprise resource