Development of Industrial Network Forensics Lessons

Thuy D. Nguyen
Naval Postgraduate School
tdnguyen@nps.edu

Cynthia E. Irvine
Naval Postgraduate School
irvine@nps.edu

ABSTRACT

Most forensic investigators are trained to recognize abusive network behavior in conventional information systems, but they may not know how to detect anomalous traffic patterns in industrial control systems (ICS) that manage critical infrastructure services. We have developed and laboratory-tested hands-on teaching material to introduce students to forensics investigation of intrusions on an industrial network. Rather than using prototypes of ICS components, our approach utilizes commercial industrial products to provide students a more realistic simulation of an ICS network. The lessons cover four different types of attacks and the corresponding post-incident network data analysis.

KEYWORDS

Industrial control system, network forensics, security education

1 INTRODUCTION

To provide real-time assessment of operating conditions, it is now becoming common practice to connect industrial control systems (ICS) to vulnerable external networks. As a result, these cyber-physical systems have become active targets of cyber attacks.

In December 2014, ICS-CERT issued an alert stating that certain control system networks in the United States had been infected by the BlackEnergy malware since 2011. In December 2015, an ICS attack crippled three regional power distribution companies in Ukraine through a series of highly synchronized operations that required extensive, long-term reconnaissance of the target operating environment [13]. ICS-CERT confirmed that a BlackEnergy 3 variant was present on the Ukrainian systems, and that it appears to have been delivered via spear phishing emails [12].

IRONGATE, an ICS malware that surfaced in June 2016 [5], includes a Stuxnet-style exploit with enhancements, i.e., sandbox evasion, and hiding its activity by recording and replaying captured data. At the time of its discovery, IRONGATE seemed to be prototype malware, but its existence confirms the threat of stealthy malware lurking in known malware and virus repositories (e.g., VirusTotal) and potentially operational ICS networks. It was introduced in 2014 but not detected until late 2015 by researchers who were looking for a different exploit.

Modern industrial network protocols have evolved from serial-based fieldbus protocols to Ethernet-based protocols that operate at the application layer of the TCP/IP networking model. The Common Industrial Protocol (CIP) [17] and Ethernet/Industrial Protocol (EtherNet/IP) [18] are two well-known Ethernet-based industrial network protocols that are used by a large number of industrial automation vendors. Rockwell Automation/Allen-Bradley (RA/AB) ControlLogix programmable logic controllers (PLCs) implement these protocols. A number of EtherNet/IP and CIP vulnerabilities in the ControlLogix PLCs have been identified since 2012 [11, 21].

Passive monitoring does not disrupt time-critical control processes and thus is well suited for use in an Operational Technology (OT) environment. Similar to analysis of conventional network traffic data in IT systems, interpreting the flow of industrial network traffic requires knowing the difference between the expected and observed behavior of each OT system in the network. Most forensic investigators are trained to recognize abusive network behavior in IT systems. However, industrial network protocols such as EtherNet/IP and CIP differ significantly from the traditional network protocols, and these investigators may lack the specific skills required to detect anomalous ICS traffic patterns.

As PLCs communicate with human operators who use human-machine interface (HMI) systems and with field devices, their network traffic contains information that can provide important clues about attack characteristics. This paper describes our first steps in developing course material to teach forensic analysis of industrial network traffic. It describes the following contributions:

(1) We have created a set of hands-on ICS network forensics lessons that can be used as a standalone ICS module in a traditional network forensics course, or be assimilated into more advanced lessons in an ICS forensics course. We focus on the EtherNet/IP protocols and ControlLogix PLC since, according to a user survey conducted by the ARC Advisory Group [1], Rockwell Automation has a large market share in North America.
(2) We confirmed a previously reported vulnerability in a ControlLogix EtherNet/IP communications module [9] and incorporated it into a lesson to expose students to ICS security issues in a real-life context.
(3) We established a corpus of EtherNet/IP traffic that we plan to make available to other researchers. This data corpus will allow others to study ControlLogix traffic without incurring the expensive equipment and development costs.

In the remaining sections, we provide basic information on industrial network protocols in §2 and review prior work in §3. We then discuss our experimental laboratory and data collection process in §4 and describe the attack scenarios and the associated forensics lessons in §5. We conclude in §6.

2 BACKGROUND

The Common Industrial Protocol (CIP) models each node in the network as a set of objects. The following definitions are taken from the CIP specification [17]. An *object* provides "an abstraction of a component within a product." A *class* is a set of objects that "are identical in form and behavior, but may contain different attribute values." An *object instance* is "the actual representation of a

This work was supported by NSF grant DUE-1140938. The views expressed in this material are those of the authors and do not reflect the official policy or position of the National Science Foundation, the Naval Postgraduate School, the Department of Defense, or the U.S. Government.

particular object within a class." Each object or class supports a set of common services (defined in Appendix A of [17]), and in certain products, a number of vendor-specific services. For example, the *Identity object* identifies the device, and its *Status attribute* (attribute ID=5) describes the current state of the entire device [17].

The Ethernet/Industrial Protocol transports CIP messages over standard Ethernet and TCP/IP protocols [18]. UDP-based implicit messaging supports time-critical operations, whereas TCP-based explicit messaging is used for normal operations. EtherNet/IP uses the same port number (44818) for both UDP and TCP connections. An EtherNet/IP encapsulation message inside a TCP/UDP packet has a 24-byte header and a command-specific data portion. There are both pre-defined commands and vendor-specific commands, and the same packet format is used for both requests and replies.

3 RELATED WORK

The need for security curricula with an emphasis on control systems security has been estabilised [2, 7, 8]. Several projects have developed suites of cybersecurity exercises, e.g., [4, 23], but these do not include lab exercises focused on ICS security.

Grandgenett et al. discuss several attacks that exploit EtherNet/IP and CIP vulnerabilities in a ControlLogix PLC [9, 10]. They describe two EtherNet/IP denial-of-service attacks, one TCP connection hoarding attack, and two CIP attacks that allow an attacker to bypass the PLC authentication mechanism and to remotely send privileged commands to the PLC.

Folkerth proposes a forensics program and presents techniques for detecting and analyzing ICS compromises using captured network data [6]. This work shows that analyzing packet trace files can help identify reconnaissance activities and targeted PLC exploitations. While some of their experiments are similar to ours, their result is not for classroom use.

NETRESEC maintains a repository of publicly available ICS packet trace files that were captured during various ICS attack challenges [16]. The PLCs used in DigitalBond S4x15 capture-the-flag competitions were Schneider Electric Modicon, Allen Bradley MicroLogix, Advantech ADAM, and Wago PLC [3]. The PLCs used in the ICS Geek Lounge lab at the 2015 4SICS ICS summit were DirectLogic 205, Siemens S7-1200, and xLogic x-Messenger [15]. In contrast, we use ControlLogix PLCs to generate our data corpus.

Last, Morris and Gao describe four datasets of industrial network traffic that were created from various attacks against two simulated SCADA systems [14]. This work closely relates to our research except that it focuses on the MODBUS protocol while our work concentrates on the EtherNet/IP protocol.

4 APPROACH

Threat Model. Any given weakness in a constituent element of a control system (redundancy mechanisms, monitoring and logging services, etc.) can adversely impact its security posture. NIST SP 800-82 discusses a number of common attacks such as reprograming PLC firmware, sending false status to the operator, modifying the functionality of safety mechanisms, manipulating configuration settings, and denying control actions [22].

We formulate the attack scenarios summarized in Table 1 and discussed in §5. These attack scenarios assume the existence of a rogue machine on the control network and a malicious insider with unfettered access to the target PLC.

Table 1: Attack scenarios

Type	Attack scenarios
Reconnaissance	Probing PLC functionalities
Data exfiltration	Exploiting services with high-value data
Malicious insider	Performing unauthorized PLC operations
Denial-of-service	Exploiting an EtherNet/IP vulnerability to silence a target PLC

Test Environment. There are three modular PLC racks and one remote I/O unit in our test enviroment. The PLCs implement the producer-consumer communications model to share data among themselves. We only use a subset of the equipment to develop the forensics lessons. This lab environment consists of a Windows HMI system, a Hirschmann industrial switch, a Linux test workstation configured as a rogue system to launch attacks on the PLC, a PLC rack, and a network monitor (Fig. 1).

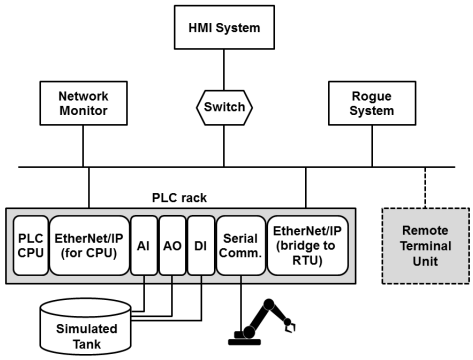

Figure 1: Lesson development laboratory

The PLC rack is an RA/AB 1756 ControlLogix system [19] that consists of a controller (CPU) module and multiple I/O modules—an EtherNet/IP communication module used to communicate with the HMI system, an analog input (AI) module, an analog output (AO) module, a digital input (DI) module, a serial communication module, and a second EtherNet/IP module used to communicate with the RTU. These I/O modules communicate with the controller module via a proprietary backplane. The controller runs a ladder logic application that controls two field devices: a simulated fluid tank system (commonly found on ships) and a robotic arm—a stand-in for any serial-connected, motorized apparatus, e.g., a rudder system. The HMI system runs the RA Studio 5000 Logix Designer software [20], which is used to develop and run ladder logic applications on the controller. The HMI system regularly asks the controller for I/O statuses and the controller's current operating mode.

Data Collection. The network packet analyzer Wireshark was used to capture and analyze network traffic between the HMI systems and the PLCs, and between the PLCs themselves. Wireshark supports a larger number of protocol display filters, including one filter for EtherNet/IP and several filters for different CIP objects.

We constructed a data corpus containing four datasets of network trace files captured by Wireshark (Table 2). The organization of the datasets is based on specific operations from the observed PLC(s).

Table 2: Summary of data corpus

ID	Description
Dataset A	Single PLC with local I/O operations
Dataset B	Single PLC with remote I/O operations
Dataset C	Multiple PLCs with producer-consumer operations
Dataset D	Single PLC with malicious operations

Datasets A, B, and C are captures of network traffic with no malicious activity. Dataset A contains EtherNet/IP packets used to control the simulated fluid tank system and the robotic arm, including the use of a simulated HMI override mechanism that allows a local operator to control the equipment in case of an emergency. The data captured in Dataset B represent the communications between a master PLC and a slave remote I/O unit to control a remote field device—a simulated steering system. One *producer PLC* and two *consummer PLCs* are used to generate the producer-consumer traffic for Dataset C. The producer PLC broadcasts a shared *tag* (variable) and multiple consumer PLCs can simultaneously receive the same tag without additional programming logic.

Dataset D is the only one with malicious activity. It consists of six packet trace files that are used to develop the forensics lessons. The lessons and packet statistics associated with the capture file used in each lesson are shown in Table 3.

Table 3: Summary of packet trace files in dataset D

ID	Forensics Lesson	Total Packets	EtherNet/IP Packets
Capture I	Reconnaissance—PLC information	14232	390 (2.7%)
Capture II	Reconnaissance—TCP/IP services	1829	312 (17.1%)
Capture III	Data exfiltration	1177	294 (25.0%)
Capture IV	Malicious insider—normal activity	90	50 (55.6%)
Capture V	Malicious insider—abnormal activity	848	490 (57.8%)
Capture VI	Denial-of-service	202006	42660 (21.1%)

5 LESSON DEVELOPMENT

The objective of our lessons is to introduce students to common vulnerabilities in an industrial network and a commercial EtherNet/IP implementation, and to demonstrate the importance of industrial network data analysis in forensics investigations. Students use Wireshark to analyze packet capture (PCAP) files in Dataset D and answer a series of questions in a written report to demonstrate their comprehension of the assigned tasks. The students must provide evidence from the PCAP files that supports their answers. A naïve way to ensure that each student performs his or her own work is to change the timestamp values in the original PCAP files to make a unique PCAP file for each student. This parameterization can easily be done via a script that uses Wireshark's `editcap` utility to make the time adjustment. There are four lessons; each corresponds to an attack scenario shown in Table 1. The malicious insider attack is against the ControlLogix CPU module whereas the other three attacks are against the Ethernet/IP module.

5.1 Lesson 1: Reconnaissance

Attack Scenario. This attack is against the integrated web server running on the EtherNet/IP module, which, when contacted, returns a list of available operations, some of which require user login with appropriate access permissions. However, it is common knowledge that ICS operators tend to use vendor-default or easy-to-remember passwords for ease of maintenance. The ControlLogix EtherNet/IP module's default password is included in SCADAPASS [1]. Unprotected operations that do not require login can be exploited to obtain information about the web server, and the configuration of the PLC rack, including the model number and firmware version of each module in the rack, and the running status of each PLC module. This information allows an attacker to construct malware targeting the specific web server and PLC modules.

Lesson Description. This lesson demonstrates how a rogue machine on the network can discover information about the PLC configuration and its embedded services from browsing the web server running on the ControlLogix EtherNet/IP module. This lesson consists of two exercises. In the first exercise, students examine the Capture I packet trace to reconstruct the browsing activities used by the attacker to learn about the different modules installed in a PLC rack. Vendor documentation (available on the Internet) indicates that the ControlLogix EtherNetIP modules support several TCP/IP application services. In the second exercise, students inspect the Capture II packet trace to determine how the attacker uncovers which services are enabled.

5.2 Lesson 2: Data Exfiltration

Attack Scenario. The ControlLogix EtherNet/IP module provides other TCP/IP services, one of which is file transfer via the File Transfer Protocol (FTP). This attack scenario focuses on the FTP server since, while inherently insecure, FTP is generally used in ICS. Similar to other FTP servers, the ControlLogix FTP server also transmits the password in plain text.

Files stored on the EtherNet/IP module are divided into two categories: normal-access and special-access. Special-access files such as module configuration files can only be accessed when the module is in backup/restore mode. Normal-access files can always be accessed and include Electronic Data Sheet (EDS) files [17] that contain configuration data such as the module's firmware version, device characteristics, and TCP/IP capability. An attacker can use the information in the EDS files to craft a targeted malware.

Lesson Description. After finding out which application services are enabled on a controller, the attacker will exploit those services

[1] A list of default passwords of ICS products maintained by the SCADA StrangeLove research group.

to obtain high-value control data. The objective of this lesson is to show how such data can be extracted if a vulnerable service such as FTP is left running on the PLC. Students analyze the retrieved data captured in the Capture III packet trace and discuss a hypothetical attack scenario based on their understanding of the EDS data.

5.3 Lesson 3: Controller Manipulation

Attack Scenario. A mechanical mode switch on the front of the ControlLogix CPU module is used to set the operating mode of the controller. The location of this switch affords an adversary a quick and non-intrusive mechanism to alter the operating behavior of the controller, which can greatly affect the industrial control process. The CPU module can be in one of five operating modes: *run*, *program*, *remote run*, *remote program*, and *remote test* [19]. An insider with physical access to the PLC rack can use the mode switch to manipulate the CPU's operating mode.

Lesson Description. The main task is to analyze CIP messages in two packet traces (Capture IV and Capture V) to determine the actions taken by the insider. This lesson utilizes the Identity object used in the Get Attribute List (GAL) and Multiple Service Packet (MSP) services. The GAL service returns the requested object attribute(s); the MSP service allows a CIP server (e.g., the controller) to perform multiple services in a single CIP message. In this lesson, the HMI system (a CIP client) uses the MSP service with multiple embedded GAL requests to obtain device and I/O information from different modules in the PLC rack. One of those GAL requests returns the current status of the CPU module via the Status attribute of the Identity object. Students must 1) interpret EtherNet//IP and CIP traffic, 2) describe how the CPU module uses the CIP Identity object to report its operating mode at a given time, and 3) compare and contrast normal traffic and abnormal traffic to detect unauthorized operating mode changes.

5.4 Lesson 4: Resource Exhaustion

Attack Scenario. This attack utilizes the RegisterSession (RS) command to cause a resource exhaustion condition on the Ethernet/IP module. This command must be executed to establish an EtherNet/IP session prior to any CIP communications between two devices. The unique 32-bit session handle returned in an RS reply is used in all subsequent EtherNet/IP messages until session termination. The RS data in the Capture VI trace indicate that the EtherNet/IP module increments the session handle value by one hundred hexadecimal (0x100). By repeatedly sending the RS command without terminating any existing sessions, we observe that the EthereNet/IP module eventually stops responding to the RS command after 17000 (approx.) requests. The CIP messages between the HMI system and the PLC also cease afterward. This attack leverages the EtherNet/IP exploit described by Grandgenett [9].

Lesson Description. This lesson shows that a deficiency in a protocol implementation could be exploited to silently bring down a PLC. The narrative is as follows: The HMI system reports that EtherNet/IP communications with the PLC has been lost, but there is no error indication at the controller module or the EtherNet/IP module, i.e., the status LEDs on the front panel of those modules still indicate that the modules are running normally. The HMI system can still *ping* the EtherNet/IP module, however subsequent requests to establish EtherNet/IP sessions continue to fail. This problem persists until the PLC system is reset. Students must describe the way the RS command is used to cause the observed denial-of-service condition, as well as the algorithm used by the ControlLogix EtherNet/IP implementation to allocate session handles.

6 CONCLUSION

This paper discusses four hands-on lessons to teach ICS network forensics using commercial ICS components. They are intended to help students gain technical proficiency in assessing and reconstructing security incidents, ranging from seemingly innocuous reconnaissance operations to the massive consumption of a critical resource that can silence a PLC.

The four lessons and our EtherNet/IP network data corpus will be publicly disseminated after student "field-testing" in a new computer science course that addresses security in cyber-physical systems. As our ongoing investigation of vulnerabilities in industrial protocol implementations continues, we plan to develop additional lessons that transfer our findings into classroom activities.

REFERENCES

[1] ARC Advisory Group. 2001. ARC User Survey: PLC Supplier Preferences. (2001).
[2] Department of Homeland Security (U.S.). 2008. Critical Infrastructure and Control Systems Security Curriculum Version 1.0. (2008).
[3] Digital Bond, Inc. 2015. ICS Village. (2015). Retrieved January 14, 2017 from http://www.digitalbond.com/s4/s4x15-week/s4x15-ics-village/
[4] Du, W., Jayaraman, K., Gaubatz, N.B. 2010. Enhancing Security Education with Hands-on Laboratory Exercises. In *5th Annual Symposium on Information Assurance (ASIA'10)*. 56–61.
[5] FireEye Inc. 2016. IRONGATE ICS Malware: Nothing to See Here...Masking Malicious Activity on SCADA Systems. (2016).
[6] Folkerth, L. 2015. Forensic Analysis of Industrial Control Systems. (2015).
[7] Foo, E., Branagan, M., Morris, T. 2013. A Proposed Australian Industrial Control System Security Curriculum. In *46th Hawaii Intl. Conf. on Sys. Sci.* . 1754–1762.
[8] Francia III, G.A. 2011. Critical Infrastructure Security Curriculum Modules. In *2011 Information Security Curriculum Development Conference*. 54–58.
[9] Grandgenett, R., Gandhi, R., and Mahoney, W. 2014. Exploitation of Allen Bradley's Implementation of EtherNet/IP for Denial of Service Against Industrial Control Systems. In *9th International Conference on Cyber Warfare and Security*.
[10] Grandgenett, R., Gandhi, R., and Mahoney, W. 2015. Authentication Bypass and Remote Escalated I/O Command Attacks. In *10th Annual Cyber and Information Security Research Conference (CISR '15)*.
[11] Industrial Control Systems CERT. 2013. Advisory (ICSA-13-011-03) Rockwell Automation ControlLogix PLC Vulnerabilities. (2013).
[12] Industrial Control Systems CERT. 2014. Alert (IR-ALERT-H-14-281-01E) Ongoing Sophisticated Malware Campaign Compromising ICS (Update E). (2014).
[13] Industrial Control Systems CERT. 2016. Alert (IR-ALERT-H-16-056-01) Cyber-Attack Against Ukrainian Critical Infrastructure. (2016).
[14] Morris, T. and Gao, W. 2014. Industrial Control System Traffic Data Sets for Intrusion Detection Research. In *Critical Infrastructure Protection VIII: 8th IFIP WG 11.10 International Conference Revised Selected Papers*. 65–78.
[15] NETRESEC AB. 2015. Capture files from 4SICS Geek Lounge. (2015). Retrieved January 14, 2017 from http://www.netresec.com/?page=PCAP4SICS
[16] NETRESEC AB. 2015. Publicly available PCAP files. (2015). Retrieved January 14, 2017 from http://www.netresec.com/?page=PcapFiles
[17] ODVA & ControlNet International Ltd. 2007. The CIP Networks Library Volume 1, Common Industrial Protocol (CIP). (2007).
[18] ODVA & ControlNet International Ltd. 2007. The CIP Networks Library Volume 2, EtherNet/IP Adaptation of CIP. (2007).
[19] Rockwell Automation 2014. *ControlLogix System User Manual.* Rockwell Automation. Publication 1756-UM001O-EN-P.
[20] Rockwell Automation 2016. *Logix5000 Controllers Ladder Diagram.* Rockwell Automation. Publication 1756-PM008F-EN-P.
[21] Santamarta, R. 2012. Attacking ControlLogix. (2012).
[22] Stouffer, K., Lightman, S., Pillitteri, V., Abrams, M., Hahn, A. 2015. Forensic Analysis of Industrial Control Systems. (2015). Revision 2.
[23] X. Wang, Y. Bai, and G. C. Hembroff. 2015. Hands-on Exercises for IT Security Education. In *Proc. 16th Ann. Conf. on Information Technology Education*. 161–166.