Poster: Improving SCADA Security with Context-aware Network Profiling

Dina Hadžiosmanović*, Robin Sommer†‡, Damiano Bolzoni* and Pieter Hartel*
*University of Twente, Enschede, The Netherlands
†International Computer Science Institute, Berkeley, CA, USA
‡Lawrence Berkeley National Laboratory, Berkeley, CA, USA

***Abstract*—SCADA (Supervisory Control and Data Acquisition) systems are computer systems used for monitoring and controlling industrial processes such as power plants and power grid systems, water, gas and oil distribution, production systems for food, cars and other products.**

In the current effort we propose a new approach for deviation monitoring in SCADA networks. We use statistics extracted from parsed protocol messages to build models of usual plant operations.

We validate our approach using network traffic from five real-life SCADA installations operating on three common industrial protocols.

INTRODUCTION

The security of SCADA systems has gained increasing attention recently. There are two main reasons for this. First, SCADA systems often control critical infrastructures. This implies that a failure of SCADA systems may endanger people health and safety, damage industrial facilities and produce financial loss. Secondly, recent reports show a significant number of security incidents in these environments. For example, a security study of 291 utility and energy companies in the U.S. [4] states that 76% of the companies report that they suffered one or more security incidents during the past 12 months.

The main reason for that is the fact that SCADA systems are not built with security in mind. Introducing security solutions into current facilities, however, is not a trivial task. This is because SCADA systems significantly differ from traditional computer systems. For example, SCADA systems typically use proprietary architectures and communication protocols. A common strategy for increasing the security in computer systems is deploying systems that monitor network traffic and alert on specific attack signatures. Such approach can only detect known attacks and thus cannot protect against unknown threats. In the SCADA context, the number of known attacks is relatively low, thus there are still few attack signatures available. This implies that the signature-based approach is even less effective in SCADA environments.

Another approach relies on describing common system operations where any behaviour significantly different from common operation implies a potential threat. SCADA systems seem to have promising properties for this approach to be effective. In particular, the behaviour of SCADA systems is less dynamic than traditional computer networks (e.g., static IP addresses, a limited number of services, dedicated devices, semi-automated procedures). Due to these reasons, we assume that SCADA are stable and thus more predictable [6]. These features are desirable when deriving models of normal behaviour.

However, profiling SCADA environment still remains a challenging task. This is due to the fact that SCADA systems run on various platforms with different implementations in place. Also, the semantic interpretation of data values sent over the network depends on the process context (e.g., current process state). This means that a proper security solution needs to understand data passed around [1].

PROBLEM

The state-of-the-art approaches for describing normal network operation do not capture context information required for interpreting a particular SCADA command. There are two reasons for this.

First, some approaches simply do not take into account the semantic meaning of command messages. For example, state-of the art approaches for monitoring network traffic are typically based on the quantitative analysis of the network activity (e.g., host profiles built on flow-based statistics) [8]. The analysis on this level monitors trends in the communication but simply cannot distinguish different types of command messages and describe different ways of operation.

Secondly, some approaches do analyse the underlying protocol but do not explore the context of the extracted command. This typically leads to defining strict policies that do not always match the context. For example, monitoring systems in [3][7] are designed to observe and specify legitimate operations over time (e.g., types of commands sent over the network). However, this implies using static, typically a manually-defined specification that is hard to maintain and might not accurately describe legitimate behaviour (e.g., a given policy is: *"workstation A can never write to the process controller"* while a desired policy is: *"workstation A can write to the process controller if the workstation is operated by user X"*).

APPROACH

We believe that adding more context to the profiling approach will make more robust profiles of normal operations. Thus we propose an approach that combines a qualitative and quantitative analysis of messages passed through the network. The qualitative analysis refers to monitoring parsed message types and parameter values, rather than coarse network statistics (such as flow-based statistics). The quantitative analysis refers to extracting statistical properties from observed message types and thus monitoring general trends of common operation. By doing both analyses, we improve the context knowledge about a particular operation.

We perform our experiments on five different real-life SCADA installations that serve as gas distribution and water purification. The analysed installations operate on three common industrial network protocols: ModbusTCP, MMS and IEC104.

We intend to implement our profiling approach in two steps.

First we extract a number of statistical properties to describe normal plant operations. In particular, we analyse various types and frequencies of messages sent over the network. Second, we analyse the system behaviour over time and estimate how predictable is the behaviour of specific devices. In this way we evaluate the trust in derived profiles.

For the first step, we leverage the Bro IDS [5]. The framework enables us to aggregate the information about the way the network protocol is used in a particular installation. We perform an extensive analysis of potentially useful features and choose a set of most promising ones.

We do this by:

- analysing known SCADA-related exploits - leverage the knowledge about vulnerabilities in various SCADA protocols (e.g., using [2]) to identify invariant properties that could distinguish normal operation from malicious action,
- perform statistical test of significance - perform statistical tests to explore the most descriptive features that define common operation,
- use expert knowledge - use the knowledge gained during the interview with SCADA users to identify interesting, and potentialy significant properties.

Some of the extracted properties are: number of messages sent in the time unit, type and subtype of messages sent, number of devices that operate specific type of message, resources used, values used in specific message (sub)types, time passed between particular message types, etc.

For the second step of our profiling approach, we leverage the fact that a SCADA systems consist of several dedicated (and semi-automated) device roles (such as PLCs, client machines, several server types) that repeat across various deployments in a consistent way. By analysing the behaviour of devices with the same role, we estimate which roles describe more predictable behaviour. Similarly, by observing statistical properties of each device over time we estimate the trust into derived profiles.

I. CONTRIBUTION

The main contributions of our work are:

- we provide the characterization of common operation in several production sites,
- we propose a new approach for deviation monitoring in SCADA networks,
- we provide the technology that implements our monitoring approach,
- we perform validation experiments using several real world SCADA installations.

REFERENCES

[1] J. Bigham, D. Gamez, and N. Lu. Safeguarding SCADA systems with anomaly detection. In *MMMACNS '03: Proc. 2nd International Workshop on Mathematical Methods, Models and Architectures for Computer Network Security*, LNCS 2776, pages 171–182. Springer Verlag, 2003.

[2] Digital Bond. *Project Basecamp*, accessed March, 2012. http://www.digitalbond.com/tools/basecamp/.

[3] Steven Cheung, Bruno Dutertre, Martin Fong, Ulf Lindqvist, Keith Skinner, and Alfonso Valdes. Using model-based intrusion detection for SCADA networks. In *Proceedings of the SCADA Security Scientific Symposium*, 2007.

[4] Ponemon Institute. State of IT Security: Study of Utilities and Energy Companies, 2011.

[5] V. Paxson. Bro: a system for detecting network intruders in real-time. *Comput. Netw.*, 31:2435–2463, December 1999.

[6] Keith Stouffer, Joe Falco, and Karen Scarfone. Guide to Industrial Control Systems (ICS) Security. NIST Special Publication 800-82, 2011.

[7] Tofino. *Tofino security applicance*, accessed March, 2012. https://www.tofinosecurity.com/products/tofino-security-appliance.

[8] Alfonso Valdes and Steven Cheung. Communication pattern anomaly detection in process control systems. In *2009 IEEE International Conference on Technologies for Homeland Security*, Waltham, MA, May 11–12, 2009.