INTRUSION DETECTION AND CORRELATION

Challenges and Solutions

Advances in Information Security

Sushil Jajodia
Consulting editor
Center for Secure Information Systems
George Mason University
Fairfax, VA 22030-4444
email: jajodia@gmu.edu

The goals of Kluwer International Series on ADVANCES IN INFORMATION SECURITY are, one, to establish the state of the art of, and set the course for future research in information security and, two, to serve as a central reference source for advanced and timely topics in information security research and development. The scope of this series includes all aspects of computer and network security and related areas such as fault tolerance and software assurance.

ADVANCES IN INFORMATION SECURITY aims to publish thorough and cohesive overviews of specific topics in information security, as well as works that are larger in scope or that contain more detailed background information than can be accommodated in shorter survey articles. The series also serves as a forum for topics that may not have reached a level of maturity to warrant a comprehensive textbook treatment.

Researchers as well as developers are encouraged to contact Professor Sushil Jajodia with ideas for books under this series.

Additional titles in the series:

THE AUSTIN PROTOCOL COMPILER by Tommy M. McGuire and Mohamed G. Gouda; ISBN: 0-387-23227-3

ECONOMICS OF INFORMATION SECURITY by L. Jean Camp and Stephen Lewis; ISBN: 1-4020-8089-1

PRIMALITY TESTING AND INTEGER FACTORIZATION IN PUBLIC KEY CRYPTOGRAPHY by Song Y. Yan; ISBN: 1-4020-7649-5

SYNCHRONIZING E-SECURITY by Godfried B. Williams; ISBN: 1-4020-7646-0

INTRUSION DETECTION IN DISTRIBUTED SYSTEMS:
An Abstraction-Based Approach by Peng Ning, Sushil Jajodia and X. Sean Wang ISBN: 1-4020-7624-X

SECURE ELECTRONIC VOTING edited by Dimitris A. Gritzalis; ISBN: 1-4020-7301-1

DISSEMINATING SECURITY UPDATES AT INTERNET SCALE by Jun Li, Peter Reiher, Gerald J. Popek; ISBN: 1-4020-7305-4

SECURE ELECTRONIC VOTING by Dimitris A. Gritzalis; ISBN: 1-4020-7301-1

APPLICATIONS OF DATA MINING IN COMPUTER SECURITY, edited by Daniel Barbará, Sushil Jajodia; ISBN: 1-4020-7054-3

MOBILE COMPUTATION WITH FUNCTIONS by Zeliha Dilsun Kırlı, ISBN: 1-4020-7024-1

Additional information about this series can be obtained from
http://www.wkap.nl/prod/s/ADIS

INTRUSION DETECTION AND CORRELATION

Challenges and Solutions

by

Christopher Kruegel
Fredrik Valeur
Giovanni Vigna

University of California, Santa Barbara, USA

Springer

eBook ISBN: 0-387-23399-7
Print ISBN: 0-387-23398-9

Created in the United States of America

Visit Springer's eBookstore at: http://ebooks.kluweronline.com
and the Springer Global Website Online at: http://www.springeronline.com

Contents

List of Figures

List of Tables

Preface

The Internet is omnipresent and companies have increasingly put critical resources online. This has given rise to the activities of cyber criminals, and virtually all organizations face increasing threats to their networks and the services they provide. This book presents intrusion detection systems (IDSs) and addresses the problem of managing and correlating the alerts that are produced. We discuss the role of intrusion detection in the realm of network security and compare it to traditional methods such as firewalls and cryptography. We then analyze the challenges in interpreting and combining (i.e., correlating) alerts produced by these systems. Existing academic and commercial systems are classified and their advantages and shortcomings are presented, especially in the case of deployment in large, real-world sites.

Recently, IDSs have been increasingly pounded for failing to meet the expectations that researchers and IDS vendors were rising. Promises that IDSs are capable of reliably identifying malicious activity in large networks were premature and never turned into reality. While virus scanners and firewalls have visible benefits and remain virtually unnoticed during normal operations, the situation is different with intrusion detection sensors. State-of-the-art IDSs produce hundreds or even thousands of alerts every day. Unfortunately, almost all of these alerts are false positives, that is, they are not related to security-relevant incidents. Although tuning and proper configuration eliminate the most obvious false alerts, the problem of the vast imbalance between important and spurious notifications remains.

Researchers and IDS vendors have reacted and proposed alert correlation, an additional step intended to manage the alert flood and turn the raw sensor output into compact reports on the security status of the network under surveillance. The idea is to aggregate and group individual alerts into attack scenarios that provide a higher-level view of the activities on the network. Unfortunately, current systems fall short in dealing with the immense data volume that is produced by the sensors that are deployed in large network installations. In

addition, dedicated nodes such as centralized processors become vulnerable to faults or targeted denial of service attempts and often represent performance bottlenecks. Another problem stems from the fact that it is often the case that sensor alerts are invalid. This causes the correlation process to deduce attack scenarios from incidents that have never occurred.

We address the aforementioned issues and present solutions that allow intrusion detection systems to be deployed in real-world installations to the benefit of the system administrator. Our proposed alert correlation process is realized by collaborating nodes that correlate and assemble the pieces of evidence, which are scattered over many hosts in the victim's network, into a single and coherent picture of ongoing attacks. The information of emerging threats is then fed back into the system and utilized to selectively adapt to data from suspicious sources. The main focus of our design is the protection of huge enterprise networks against coordinated attacks without being overwhelmed by the produced alert data and without failing because of the loss of a few critical correlation nodes. We also describe an approach to reduce the number of false positives by actively performing alert verification. The idea is to determine whether a potential attack has succeeded by checking for visible traces that this attack has left' on the system.

This book introduces solutions to practical problems that intrusion detection systems experience when deployed in large network installations. The reader is familiarized with the basics and concepts of this fast growing and fascinating field in network security and learns about state-of-the-art systems. We focus on current research problems and help the reader understand the limitations and advantages of intrusion detection systems and, in particular, alert correlation and mechanisms to detect false alarms.