Available online at www.sciencedirect.com
SCIENCE DIRECT®
Electronic Notes in Theoretical Computer Science
ELSEVIER
Electronic Notes in Theoretical Computer Science 125 (2005) 3–12
www.elsevier.com/locate/entcs

Believing the Integrity of a System
(Invited Talk)

Simon N. Foley[1,2]

Department of Computer Science
University College Cork
Ireland

Abstract

An integrity policy defines the situations when modification of information is authorised and is enforced by the protection mechanisms of a system. Traditional models of protection tend to define integrity in terms of ad-hoc authorisation techniques whose effectiveness are justified more on the basis of experience and "best practice" rather than on any theoretical foundation. In a complex application system it is possible that an integrity policy may have been incorrectly configured, or that the protection mechanisms are inadequate, resulting in an unexpected system compromise. This paper examines the meaning of integrity and and describes a simple belief logic approach for analysing the integrity of a system configuration.

Keywords: Security, Integrity, Security Protocols, Belief logics, System Configuration.

1 Introduction

The 2001 Computer Crime and Security Survey [9] reported that 196 of the respondents to the survey could quantify their losses due to unauthorised use of computer systems at a total of US$378 million in the previous year. While access-control mechanisms, firewalls and so forth may help counter such losses, we can never be confident about security unless we are provided with some assurance of their effectiveness. Such assurance may be achieved, in part, by analysing whether a formal description of the system upholds certain security

[1] Supported by Enterprise Ireland Basic Research Grant Scheme SC/2003/7/ *FITNYSS Foundations for Integrity and Analysis.*
[2] Email: s.foley@cs.ucc.ie

1571-0661/$ – see front matter
doi:10.1016/j.entcs.2004.09.037

properties. These properties include confidentiality (no unauthorised release of information) and integrity (no unauthorised modification of information). The study of integrity as a formal security property has received little attention within the research community; confidentiality has been extensively studied and is the better understood of the two properties.

Early security research [3] characterised integrity in terms of read-write access controls between subjects and objects. This provides for a very coarse interpretation of integrity [30]; for example, once granted access to an account database, a bank clerk can make any change to the customer's account details. Access triples, well-formed transactions, and the principles of encapsulation [8,25], provide finer grained control by constraining the operations that a subject may carry out on an object: the bank clerk may execute only deposit or withdraw operations to access an account database.

Many integrity compromises are a result of 'insiders' executing fraudulent but authorised operations [9]. For example, the bank clerk executes an account deposit without lodging actual funds. Separation of duty [8,13,37,35] controls decrease the potential for fraud by involving at least two individuals at different points in a transaction: for example, by reconciling bank accounts and funds received each day, a supervisor detects and corrects the fraudulent deposit by the clerk. Role Based Access Control models [31,32] and authorisation models [2,21] provide integrity controls based on structures that organise related operations into roles and constrain the way that roles may be assigned and/or inherited by users; separation of duty is expressed within these models using role constraints.

These conventional security models describe *controls* for achieving integrity; they take an operational and/or implementation oriented approach by defining *how* to achieve integrity. No attempt is made to formalise a *property* that defines *what* is meant by integrity. For example, [8] recommends a combination of separation of duties, access-triples and auditing as a strategy for achieving integrity: it does it not attempt to address what is meant by integrity. Confidence is achieved to the extent that good design principles have been applied; there is no assurance that an integrity property is upheld. Thus, when we define a complex separation of duty policy we do not know, for certain, whether a dishonest user can bypass the intent of the separation via some unexpected circuitous, but authorised, route.

Jacob [20] formalises integrity as a functional property. This interpretation of integrity means that an integrity mechanism determines whether the current request for an operation is authorised based on the history of past authorisation requests that led to the current state.

In our research [14,15], we argue that integrity should be regarded as a