A Comparison of Commercial and Military Computer Security Policies

David D. Clark* - David R. Wilson**

* Senior Research Scientist, MIT Laboratory for Computer Science
545 Technology Square, Cambridge, MA 02139
** Director, Information Security Services, Ernst & Whinney
2000 National City Center, Cleveland, OH 44114

ABSTRACT

Most discussions of computer security focus on control of disclosure. In particular, the U.S. Department of Defense has developed a set of criteria for computer mechanisms to provide control of classified information. However, for that core of data processing concerned with business operation and control of assets, the primary security concern is data integrity. This paper presents a policy for data integrity based on commercial data processing practices, and compares the mechanisms needed for this policy with the mechanisms needed to enforce the lattice model for information security. We argue that a lattice model is not sufficient to characterize integrity policies, and that distinct mechanisms are needed to control disclosure and to provide integrity.

INTRODUCTION

Any discussion of mechanisms to enforce computer security must involve a particular security policy that specifies the security goals the system must meet and the threats it must resist. For example, the high-level security goals most often specified are that the system should prevent unauthorized disclosure or theft of information, should prevent unauthorized modification of information, and should prevent denial of service. Traditional threats that must be countered are system penetration by unauthorized persons, unauthorized actions by authorized persons, and abuse of special privileges by systems programmers and facility operators. These threats may be intentional or accidental.

Imprecise or conflicting assumptions about desired policies often confuse discussions of computer security mechanisms. In particular, in comparing commercial and military systems, a misunderstanding about the underlying policies the two are trying to enforce often leads to difficulty in understanding the motivation for certain mechanisms that have been developed and espoused by one group or the other. This paper discusses the military security policy, presents a security policy valid in many commercial situations, and then compares the two policies to reveal important differences between them.

The military security policy we are referring to is a set of policies that regulate the control of classified information within the government. This well-understood, high-level information security policy is that all classified information shall be protected from unauthorized disclosure or declassification. Mechanisms used to enforce this policy include the mandatory labeling of all documents with their classification level, and the assigning of user access categories based on the investigation (or "clearing") of all persons permitted to use this information. During the last 15 to 20 years, considerable effort has gone into determining which mechanisms should be used to enforce this policy within a computer. Mechanisms such as identification and authorization of users, generation of audit information, and association of access control labels with all information objects are well understood. This policy is defined in the Department of Defense Trusted Computer System Evaluation Criteria [DOD], often called the "Orange Book" from the color of its cover. It articulates a standard for maintaining confidentiality of information and is, for the purposes of our paper, the "military" information security policy. The term "military" is perhaps not the most descriptive characterization of this policy; it is relevant to any situation in which access rules for sensitive material must be enforced. We use the term "military" as a concise tag which at least captures the origin of the policy.

In the commercial environment, preventing disclosure is often important, but preventing unauthorized data modification is usually paramount. In particular, for that core of commercial data processing that relates to management and accounting for assets, preventing fraud and error is the primary goal. This goal is addressed by enforcing the integrity rather than the privacy of the information. For this reason, the policy we will concern ourselves with is one that addresses integrity rather than disclosure. We will call this a commercial policy, in contrast to the military information security policy. We are not suggesting that integrity plays no role in military concerns. However, to the extent that the Orange Book is the articulation of the military information security policy, there is a clear difference of emphasis in the military and commercial worlds.

While the accounting principles that are the basis of fraud and error control are well known, there is yet no Orange Book for the commercial sector that articulates how these policies are to be implemented in the context of a computer system. This makes it difficult to answer the question of whether the mechanisms designed to enforce military information security policies also apply to enforcing commercial integrity policies. It would be very nice if the same mechanisms could meet both goals, thus enabling the commercial and military worlds to share the development costs of the necessary mechanisms. However, we will argue that two distinct classes of mechanism will be required, because some of the mechanisms needed to enforce disclosure controls and integrity controls are very different.

Therefore, the goal of this paper is to defend two conclusions. First, there is a distinct set of security policies, related to integrity rather than disclosure, which are often of highest priority in the commercial data processing environment. Second, some separate mechanisms are required for enforcement of these policies, disjoint from those of the Orange Book.

MILITARY SECURITY POLICY

The policies associated with the management of classified information, and the mechanisms used to enforce these policies, are carefully defined and well understood within the military. However, these mechanisms are not necessarily well understood in the commercial world, which normally does not have such a complex requirement for control of unauthorized disclosure. Because the military security model provides a good starting point, we begin with a brief summary of computer security in the context of classified information control.

The top-level goal for the control of classified information is very simple: classified information must not be disclosed to unauthorized individuals. At first glance, it appears the correct mechanism to enforce this policy is a control over which individuals can read which data items. This mechanism, while certainly needed, is much too simplistic to solve the entire problem of unauthorized information release. In particular, enforcing this policy requires a mechanism to control writing of data as well as reading it. Because the control of writing data is superficially associated with ensuring integrity rather than preventing theft, and the classification policy concerns the control of theft, confusion has arisen about the fact that the military mechanism includes strong controls over who can write which data.

Informally, the line of reasoning that leads to this mechanism is as follows. To enforce this policy, the system must protect itself from the authorized user as well as the unauthorized user. There are a number of ways for the authorized user to declassify information. He can do so as a result of a mistake, as a deliberate illegal action, or because he invokes a program on his behalf that, without his knowledge, declassifies data as a malicious side effect of its execution.

This class of program, sometimes called a "Trojan Horse" program, has received much attention within the military. To understand how to control this class of problem in the computer, consider how a document can be declassified in a noncomputerized context. The simple technique involves copying the document, removing the classification labels from the document with a pair of scissors, and then making another copy that does not have the classification labels. This second copy, which physically appears to be unclassified, can then be carried past security guards who are responsible for controlling the theft of classified documents. Declassification occurs by copying.

To prevent this in a computer system, it is necessary to control the ability of an authorized user to copy a data item. In particular, once a computation has read a data item of a certain security level, the system must ensure that any data items written by that computation have a security label at least as restrictive as the label of the item previously read. It is this

mandatory check of the security level of all data items whenever they are written that enforces the high level security policy.

An important component of this mechanism is that checking the security level on all reads and writes is mandatory and enforced by the system, as opposed to being at the discretion of the individual user or application. In a typical time sharing system not intended for multilevel secure operation, the individual responsible for a piece of data determines who may read or write that data. Such discretionary controls are not sufficient to enforce the military security rules because, as suggested above, the authorized user (or programs running on his behalf) cannot be trusted to enforce the rules properly. The mandatory controls of the system constrain the individual user so that any action he takes is guaranteed to conform to the security policy. Most systems intended for military security provide traditional discretionary control in addition to the mandatory classification checking to support what is informally called "need to know." By this mechanism, it is possible for the user to further restrict the accessibility of his data, but it is not possible to increase the scope in a manner inconsistent with the classification levels.

In 1983, the U.S. Department of Defense produced the Orange Book, which attempts to organize and document mechanisms that should be found in a computer system designed to enforce the military security policies. This document stresses the importance of mandatory controls if effective enforcement of a policy is to be achieved within a system. To enforce the particular policy of the Orange Book, the mandatory controls relate to data labels and user access categories. Systems in division C have no requirement for mandatory controls, while systems in divisions A and B specifically have these mandatory maintenance and checking controls for labels and user rights. (Systems in Division A are distinguished from those in B, not by additional function, but by having been designed to permit formal verification of the security principles of the system.)

Several security systems used in the commercial environment, specifically RACF, ACF/2, and CA-TopSecret, were recently evaluated using the Orange Book criteria. The C ratings that these security packages received would indicated that they did not meet the mandatory requirements of the security model as described in the Orange Book. Yet, these packages are used commonly in industry and viewed as being rather effective in their meeting of industry requirements. This would suggest that industry views security requirements somewhat differently than the security policy described in the Orange Book. The next section of the paper begins a discussion of this industry view.

COMMERCIAL SECURITY POLICY FOR INTEGRITY

Clearly, control of confidential information is important in both the commercial and military environments. However, a major goal of commercial data processing, often the most important goal, is to ensure integrity of data to prevent fraud and errors. No user of the system, even if authorized, may be permitted to modify data items in such a way that assets or accounting records of the company are lost or corrupted. Some mechanisms in the system, such as user authentication, are an integral part of enforcing both the commercial and military policies. However, other mechanisms are very different.

The high-level mechanisms used to enforce commercial security policies related to data integrity were derived long before computer systems came into existence. Essentially, there are two mechanisms at the heart of fraud and error control: the well-formed transaction, and separation of duty among employees.

The concept of the well-formed transaction is that a user should not manipulate data arbitrarily, but only in constrained ways that preserve or ensure the integrity of the data. A very common mechanism in well-formed transactions is to record all data modifications in a log so that actions can be audited later. (Before the computer, bookkeepers were instructed to write in ink, and to make correcting entries rather than erase in case of error. In this way the books themselves, being write-only, became the log, and any evidence of erasure was indication of fraud.)

Perhaps the most formally structured example of well-formed transactions occurs in accounting systems, which model their transactions on the principles of double entry bookkeeping. Double entry bookkeeping ensures the internal consistency of the system's data items by requiring that any modification of the books comprises two parts, which account for or balance each other. For example, if a check is to be written (which implies an entry in the cash account) there must be a matching entry on the accounts payable account. If an entry is not performed properly, so that the parts do not match, this can