Chapter 4
Self-Protecting Access Control: On Mitigating Privacy Violations with Fault Tolerance

Anne V. D. M. Kayem
University of Cape Town, South Africa

Patrick Martin
Queen's University, Canada

Selim G. Akl
Queen's University, Canada

ABSTRACT

Self-protecting access control mechanisms can be described as an approach to enforcing security in a manner that automatically protects against violations of access control rules. In this chapter, we present a comparative analysis of standard Cryptographic Access Control (CAC) schemes in relation to privacy enforcement on the Web. We postulate that to mitigate privacy violations, self-protecting CAC mechanisms need to be supported by fault-tolerance. As an example of how one might to do this, we present two solutions that are inspired by the autonomic computing paradigm[1]. Our solutions are centered on how CAC schemes can be extended to protect against privacy violations that might arise from key updates and collusion attacks.

INTRODUCTION

The ability to execute multiple transactions across a myriad of applications has made the Internet a prime platform for building Web applications. Applications like Facebook (Facebook, 2010) and MySpace (MySpace, 2010), attest to this popularity and have been rated as being the most popular social networking applications in the English speaking world. Increasingly, business organizations are taking advantage of these social networking applications and other web applications to collect personal information about consumers and likewise consumers have shown a keenness for the web as a medium of communication because of the interactivity and fast

DOI: 10.4018/978-1-61350-501-4.ch004

response time it offers. Yet, the same qualities of flexibility and interactivity that the web is famous for have become an impediment in the face of the growing incidences of data privacy violations. For example, in October 2010 a Wall Street Journal Investigation revealed that many popular Facebook applications were transmitting consumer personal information to advertising and Internet tracking companies (Slattery, 2010), (Foremski, 2010). Cases like this have fueled growing concerns, on the part of consumers, that their data can be leaked without their consent to third parties. In this section, we discuss the context in which data privacy violations occur and why this happens in spite of the fact that access control mechanisms can be implemented to protect the information.

Context and Motivation

The Internet is built on the assumption that the users of the network can be trusted to behave honestly and so do not use the system or behave in ways that could compromise the performance and/or credibility of the system. Yet, this quality of open access makes web-applications inherently vulnerable to violations of information privacy rules (accidental or intentional) that can compromise the levels of data protection that these applications promise users (Harrison, February 2007), (Sandhu, 2005), (Tanenbaum & Steen, 2007). In many cases, privacy violations occur because consumers assume that they have "correctly" applied some access control mechanism that will prevent illegal access to their information. For instance, in social networking applications, it often is the case that a user will post "confidential" information and forget to set the parameter to prevent transitive disclosures to friends of the user's friends.

As well, consumers have a tendency to naively assume that business organizations will do what they promise, while business organizations are sometimes unaware of the far-reaching consequences of certain management decisions. In the case of Facebook, one could imagine that a third-party made contact by indicating that they would like to test the popularity of a new application. Facebook probably agreed because usage of the application might attract new members. However, the acceptance agreement might not have indicated clearly that the application could not collect information about the users who choose to use the application and/or people whom the users know might be interested in using the application (Fung, Wang, Chen, & Yu, 2010). Data privacy leaks like the one we describe are a growing concern for organizations because they result in a loss of revue (Foremski, 2010).

Until recently, organizations simply focused on defining a security domain and security policies were used to control access to information. The assumption was that if correctly specified, failure (either deliberate or not), on the part of users, to adhere to data privacy policies would be unlikely. However, the emergence of concepts like service-oriented architectures and cloud computing have dissolved inter-organization boundaries. Consequently, web applications and/or services can interact flexibly across multiple security domains and in ways that are not easy to predict at runtime. Therefore, security policies and access control schemes need to be modeled or extended to cope with situations in which changes in security requirements result in privacy violations.

In this chapter, we discuss the growing need to extend access control models to enforce privacy in scenarios involving changing security requirements like the Web. More specifically, we consider the literature on the more popular access control models like mandatory, discretionary, and role-based access control, and discuss some of the ways in which these models have been extended to enforce data privacy requirements. In recent years, cryptographic access control (CAC) is has received increased attention as a method of enforcing data privacy on the Web. CAC schemes have the advantage of providing protection for data in untrustworthy environments like the Web.