

Database Security and Statistical Database Security

Edgar R. Weippl

Secure Business Austria, Austria

INTRODUCTION

In this article we will present an introduction to issues relevant to database security and statistical database security. We will briefly cover various security models, elaborate on how data analysis in data warehouses (DWH) might compromise an individual's privacy, and explain which safeguards can be used to prevent attacks.

In most companies, databases are an essential part of IT infrastructure since they store critical business data. In the last two decades, databases have been used to process increasing amounts of transactional data, such as, a complete account of a person's purchases from a retailer or connection data from calls made on a cell phone.

As soon as this data became available from transactional databases and online transactional processing (OLTP) became well established, the next logical step was to use the knowledge contained in the vast amounts of data. Today, data warehouses (DWH) store aggregated data in an optimal way to serve queries related to business analysis.

In recent years, most people have begun to focus their attention on security. Early OLTP applications were mainly concerned with integrity of data during transactions; today privacy and secrecy are more important as databases store an increasing amount of information about individuals, and data from different systems can be aggregated. Thuraisingham (2002) summarizes the requirements briefly as "*However, we do not want the information to be used in an incorrect manner.*"

All security requirements stem from one of three basic requirements: confidentiality (aka secrecy), integrity, and availability (CIA). Confidentiality refers to the requirement that only authorized subjects, that is, people or processes should be permitted to read data. Integrity means that unauthorized modifications must not be permitted. This includes both modifications by unauthorized people and incorrect modification by authorized users. To correctly perform the services

requested, the system needs to remain available; a denial-of-service compromises the requirement of availability.

Other security requirements may include privacy, non-repudiation, and separation of duties. These requirements are, however, composite requirements that can be traced back to one of the three basic requirements. Privacy, for instance, is the non-disclosure (=confidentiality) of personal data; non-repudiation refers to the integrity of transaction logs and integrity of origin. Throughout this article we will focus only on technical attacks and safeguards and not on social engineering. Social engineering is often the easiest and, in many cases, a very successful attack vector. For an in-depth coverage of social engineering we recommend (Böck, 2007).

In Section 2 we cover the most relevant access control models; in Section 3 we provide an overview of security in statistical databases. Finally, in Section 4 we highlight the essentials of securing not only the transactional and the statistical databases but the entire system.

BACKGROUND

Access Control is the most important technique or mechanism for implementing the requirements of confidentiality and integrity. Since databases were among the first large-scale systems in military applications, there is a long history of security models, dating back to the 1960s. The basic principle in all access control models is that a *subject* is or is not permitted to perform a certain *operation* on an *object*. This process is described by the triplet (s, op, o). A security policy specifies who is authorized to do what. A security mechanism allows enforcement of a chosen security policy.

One can distinguish between two fundamentally different access control mechanisms: discretionary access control (DAC) and mandatory access control (MAC). In DAC models the user decides which subject is able

to access which object to perform a certain operation. In contrast, when using MAC, the system decides who is allowed to access which resource and the individual user has no discretion to decide on access rights.

Discretionary Access Control (DAC)

In relational database management systems (DBMS), the objects that need to be protected are tables and views. Modern DBMS allow a fine granularity of access control so that access to individual fields of a record can be controlled.

By default, a subject has no access. Subjects may then be *granted* access, which can be *revoked* anytime. In most systems the creator of a table or a view is automatically granted all privileges related to it. The DBMS keeps track of who subsequently gains and loses privileges, and ensures that only requests from subjects who have the required privileges—at the time the request is executed—are allowed.

Mandatory Access Control (MAC)

Mandatory Access Control is based on system-wide policies that cannot be changed by individual users. Each object in the database is automatically assigned a security class based on the access privileges of the user who created the object.

The most widely known implementation of a MAC system is a multi-level security (MLS) system. MLS systems were first described by Bell LaPadula (Bell, 1975) in the 1960s. Each subject, which could either be a user or user program, is assigned a *clearance* for a security class. Objects are assigned security *levels*. Security levels and clearances can be freely defined as long as all items are comparable pair-wise. Most common are security classes (i.e., levels and clearances), such as, top secret (TS), secret (S), confidential (C), and unclassified (U).

Rules based on security levels and clearances govern who can read or write which objects. Today, there are only a few commercially available systems that support MAC, such as, SELinux or also Oracle DBMS (Version 9 and higher) when the Oracle Label Security (OLS) option is installed.

The main reason to use a MAC system is that it prevents inherent flaws of discretionary access control, which are commonly referred to as the Trojan horse problem. The user Alice creates a program and gives

Bob INSERT privileges for the table mySecret. Bob knows nothing about this. Alice modifies the code of an executable that Bob uses so that it additionally writes Bob's secret data to the table mySecret. Now, Alice can see Bob's secret data. While the modification of the application code is beyond the DBMS' control, it can still prevent the use of the database as a channel for secret information.

ACCESS CONTROL FOR RELATIONAL DATABASES

Role-Based Access Control (RBAC)

With RBAC (Sandhu, 2000), system administrators create roles according to the job functions defined in a company; they grant permissions to those roles and subsequently assign users to the roles on the basis of their specific job responsibilities and qualifications. Thus, roles define the authority of users, the competence that users have, and the trust that the company gives to the user. Roles define both, the specific individuals allowed to access objects and the extent to which or the mode in which they are allowed to access the object (see Sandhu & Coyne & Feinstein & Youman, 1996). Access decisions are based on the roles a user has activated (Sandhu & Ferraiolo & Kuhn, 2000).

The basic RBAC model consists of four entities: users, roles, permissions, and sessions. A user is a subject who accesses different, protected objects. A role is a named job function that describes the authority, trust, responsibility, and competence of a role member. A permission is an approval for a particular type of access to one or more objects. Permissions describe which actions can be performed on a protected object and may apply to one or more objects. Both permissions and users are assigned to roles. These assignments, in turn, define the scope of access rights a user has with respect to an object. By definition, the user assignment and permission assignment relations are many-to-many relationships.

Users establish sessions during which they may activate a subset of the roles they belong to. A session maps one user to many possible roles, which results in the fact that multiple roles can be activated simultaneously and every session is assigned with a single user. Moreover, a user might have multiple sessions opened simultaneously. Belonging to several roles, a user can

5 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/database-security-statistical-database-security/10884

Related Content

Web Page Extension of Data Warehouses

Anthony Scime (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition* (pp. 2090-2095). www.irma-international.org/chapter/web-page-extension-data-warehouses/11108

Utilizing Fuzzy Decision Trees in Decision Making

Malcolm J. Beynonm (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition* (pp. 2024-2030). www.irma-international.org/chapter/utilizing-fuzzy-decision-trees-decision/11097

Distributed Data Aggregation Technology for Real-Time DDoS Attacks Detection

Yu Chenand Wei-Shinn Ku (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition* (pp. 701-708). www.irma-international.org/chapter/distributed-data-aggregation-technology-real/10897

Promoting Critical Thinking Disposition Through Virtual Reality Serious Games

Su Jiayuanand Jingru Zhang (2024). *Embracing Cutting-Edge Technology in Modern Educational Settings* (pp. 93-118). www.irma-international.org/chapter/promoting-critical-thinking-disposition-through-virtual-reality-serious-games/336192

Privacy-Preserving Data Mining

Stanley R.M. Oliveira (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition* (pp. 1582-1588). www.irma-international.org/chapter/privacy-preserving-data-mining/11030