

---

**Access Control  
for Databases:  
Concepts and  
Systems**

---

# Access Control for Databases: Concepts and Systems

---

**Elisa Bertino**

*Purdue University  
West Lafayette, IN, USA  
bertino@cs.purdue.edu*

**Gabriel Ghinita**

*Purdue University  
West Lafayette, IN, USA  
gghinita@cs.purdue.edu*

**Ashish Kamra**

*Purdue University  
West Lafayette, IN, USA  
akamra@purdue.edu*

**now**

the essence of knowledge

Boston – Delft

## Foundations and Trends<sup>®</sup> in Databases

*Published, sold and distributed by:*

now Publishers Inc.  
PO Box 1024  
Hanover, MA 02339  
USA  
Tel. +1-781-985-4510  
[www.nowpublishers.com](http://www.nowpublishers.com)  
[sales@nowpublishers.com](mailto:sales@nowpublishers.com)

*Outside North America:*

now Publishers Inc.  
PO Box 179  
2600 AD Delft  
The Netherlands  
Tel. +31-6-51115274

The preferred citation for this publication is E. Bertino, G. Ghinita and A. Kamra, Access Control for Databases: Concepts and Systems, Foundation and Trends<sup>®</sup> in Databases, vol 3, nos 1–2, pp 1–148, 2010

ISBN: 978-1-60198-416-6

© 2011 E. Bertino, G. Ghinita and A. Kamra

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: [www.copyright.com](http://www.copyright.com)

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1-781-871-0245; [www.nowpublishers.com](http://www.nowpublishers.com); [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, [www.nowpublishers.com](http://www.nowpublishers.com); e-mail: [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

**Foundations and Trends<sup>®</sup> in  
Databases**  
Volume 3 Issues 1–2, 2010  
**Editorial Board**

**Editor-in-Chief:**

**Joseph M. Hellerstein**

*Computer Science Division*

*University of California, Berkeley*

*Berkeley, CA*

*USA*

*hellerstein@cs.berkeley.edu*

**Editors**

Anastasia Ailamaki (EPFL)

Michael Carey (UC Irvine)

Surajit Chaudhuri (Microsoft Research)

Ronald Fagin (IBM Research)

Minos Garofalakis (Yahoo! Research)

Johannes Gehrke (Cornell University)

Alon Halevy (Google)

Jeffrey Naughton (University of Wisconsin)

Christopher Olston (Yahoo! Research)

Jignesh Patel (University of Michigan)

Raghu Ramakrishnan (Yahoo! Research)

Gerhard Weikum (Max-Planck Institute)

## Editorial Scope

**Foundations and Trends<sup>®</sup> in Databases** covers a breadth of topics relating to the management of large volumes of data. The journal targets the full scope of issues in data management, from theoretical foundations, to languages and modeling, to algorithms, system architecture, and applications. The list of topics below illustrates some of the intended coverage, though it is by no means exhaustive:

- Data Models and Query Languages
- Query Processing and Optimization
- Storage, Access Methods, and Indexing
- Transaction Management, Concurrency Control and Recovery
- Deductive Databases
- Parallel and Distributed Database Systems
- Database Design and Tuning
- Metadata Management
- Object Management
- Trigger Processing and Active Databases
- Data Mining and OLAP
- Approximate and Interactive Query Processing
- Data Warehousing
- Adaptive Query Processing
- Data Stream Management
- Search and Query Integration
- XML and Semi-Structured Data
- Web Services and Middleware
- Data Integration and Exchange
- Private and Secure Data Management
- Peer-to-Peer, Sensornet and Mobile Data Management
- Scientific and Spatial Data Management
- Data Brokering and Publish/Subscribe
- Data Cleaning and Information Extraction
- Probabilistic Data Management

### Information for Librarians

Foundations and Trends<sup>®</sup> in Databases, 2010, Volume 3, 4 issues. ISSN paper version 1931-7883. ISSN online version 1931-7891. Also available as a combined paper and online subscription.

Foundations and Trends<sup>®</sup> in  
Databases  
Vol. 3, Nos. 1–2 (2010) 1–148  
© 2011 E. Bertino, G. Ghinita and A. Kamra  
DOI: 10.1561/19000000014



## Access Control for Databases: Concepts and Systems

Elisa Bertino<sup>1</sup>, Gabriel Ghinita<sup>2</sup>  
and Ashish Kamra<sup>3</sup>

<sup>1</sup> *CS Department, Purdue University, West Lafayette, IN, 47907, USA,  
bertino@cs.purdue.edu*

<sup>2</sup> *CS Department, Purdue University, West Lafayette, IN, 47907, USA,  
gghinita@cs.purdue.edu*

<sup>3</sup> *ECE Department, Purdue University, West Lafayette, IN, 47907, USA,  
akamra@purdue.edu*

### Abstract

As organizations depend on, possibly distributed, information systems for operational, decisional and strategic activities, they are vulnerable to security breaches leading to data theft and unauthorized disclosures even as they gain productivity and efficiency advantages. Though several techniques, such as encryption and digital signatures, are available to protect data when transmitted across sites, a truly comprehensive approach for data protection must include mechanisms for enforcing access control policies based on data contents, subject qualifications and characteristics, and other relevant contextual information, such as time. It is well understood today that the semantics of data must be taken into account in order to specify effective access control policies. To address such requirements, over the years the database security research community has developed a number of access control techniques and

mechanisms that are specific to database systems. In this monograph, we present a comprehensive state of the art about models, systems and approaches proposed for specifying and enforcing access control policies in database management systems. In addition to surveying the foundational work in the area of access control for database systems, we present extensive case studies covering advanced features of current database management systems, such as the support for fine-grained and context-based access control, the support for mandatory access control, and approaches for protecting the data from insider threats. The monograph also covers novel approaches, based on cryptographic techniques, to enforce access control and surveys access control models for object-databases and XML data. For the reader not familiar with basic notions concerning access control and cryptography, we include a tutorial presentation on these notions. Finally, the monograph concludes with a discussion on current challenges for database access control and security, and preliminary approaches addressing some of these challenges.

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	An Historical Perspective	4
1.2	Recent Research Directions	7
1.3	Organization of the Monograph	9
<b>2</b>	<b>Background</b>	<b>11</b>
2.1	Access Control Models	11
2.2	Cryptographic Preliminaries	32
2.3	Summary	35
<b>3</b>	<b>Foundations of Access Control for Relational Database Systems</b>	<b>37</b>
3.1	The System R Access Control Model	37
3.2	Content-based Access Control	42
3.3	Mandatory Access Control Models	45
3.4	Summary	49
<b>4</b>	<b>Case Studies</b>	<b>51</b>
4.1	SQL Server 2008	51
4.2	Oracle Virtual Private Database	58
4.3	Labeled Oracle	62
4.4	Summary	66



<b>5 Fine-Grained Access Control Models and Mechanisms</b>	<b>67</b>
5.1 Fine Grained Access Control through Query Rewriting	70
5.2 SQL Language Extensions for Fine Grained Access Control	77
5.3 Fine Grained Access Control with Authorization Views	82
5.4 Summary	84
<b>6 PSAC: A Privilege State Based Access Control System</b>	<b>85</b>
6.1 Motivation	85
6.2 Design and Implementation	87
6.3 Summary	99
<b>7 Protection from Insider Threats and Separation of Duties</b>	<b>101</b>
7.1 Oracle Database Vault	102
7.2 Joint-Threshold Administration	104
7.3 Summary	111
<b>8 Access Control for Object Databases, XML Data and Novel Applications</b>	<b>113</b>
8.1 Requirements	114
8.2 The Orion Authorization Model	115
8.3 MAC Models for Object Databases	122
8.4 Access Control Models for XML Data	122
8.5 Access Control Models for Geographical Data	124
8.6 Access Control Models for Digital Libraries	126
8.7 Summary	127
<b>9 Encryption-based Access Control</b>	<b>129</b>
9.1 Encryption-based Access Control for XML Documents	130

9.2 Privacy-preserving Access Control Mechanisms	136
9.3 Summary	141
<b>10 Concluding Remarks and Research Directions</b>	<b>143</b>
<b>References</b>	<b>147</b>

# 1

---

## Introduction

---

Today all organizations rely on database systems as the key data management technology for a large variety of tasks, ranging from day-to-day operations to critical decision making. Such widespread use of database systems implies that security breaches to these systems affect not only a single user or application, but also may have disastrous consequences on the entire organization. The recent rapid proliferation of Web-based applications and information systems, and recent trends such as cloud computing and outsourced data management, has further increased the exposure of database systems and, thus, data protection is more crucial than ever. Conventional perimeter-oriented defenses, like firewalls, are inadequate in today's interconnected world and are unable to offer the fine-grained protection required for selective and secure data sharing among multiple users and applications. Security techniques offered by operating systems may offer some protection at the file system level; however the protected objects are typically files and directories and these protection units are too coarse with respect to the *logical protection units*, such as records, that are required in database systems. It is also important to appreciate that data need to be protected not only from external threats, but also from insider threats [19].

## 2 Introduction

As discussed by Bertino and Sandhu [19], data security breaches are typically classified as *unauthorized data observation*, *improper data modification*, and *data unavailability*. Unauthorized data observation results in the disclosure of information to subjects<sup>1</sup> not entitled to gain access to the information. All organizations, ranging from governmental and military organizations to social and commercial organizations, may suffer losses from both financial and human points of view as a consequence of unauthorized data observation. The unauthorized disclosure of personally identifiable data may result in privacy breaches, that may lead to identity theft and other serious consequences for the individuals. Improper data modifications, either intentional or unintentional, result in incorrect data. Any use of incorrect data may also result in heavy losses for organizations. When data are unavailable, information crucial for the proper functioning of an organization is not readily available when needed. Thus, a complete solution to data protection must meet three key requirements: (1) **secrecy** or **confidentiality** — it refers to the protection of data against unauthorized disclosures; (2) **integrity** — it refers to the prevention of improper data modifications; and (3) **availability** — it refers to the prevention and recovery from hardware and software errors and from malicious data access denials making the database system unavailable. These three requirements arise practically in all applications. Consider a database storing medical information about patients of a hospital. It is important that patient records not be released to unauthorized subjects, that records be modified only by the subjects who are properly authorized and their accuracy be assured, and that patient records be readily available to doctors in charge especially in emergency situations.

Securing data is a challenging task. It is ensured collectively by various components of a database management system (DBMS) and may also require components external to the DBMS, such as secure co-processors [1].

A key component for assuring data protection is represented by the *access control mechanism*. When a subject attempts to access some

---

<sup>1</sup>The term 'subject' refers to any active entity which tries to access the protected resources in a system. A subject can be an end-user, a process, or an application program, or an organizational role.

data, the access control mechanism checks whether or not the subject has the authorization to perform the action on the data. Authorizations are granted to subjects according to the access control policies of the organization. Confidentiality can be further enhanced by the use of encryption techniques, applied to data when being stored on secondary storage or transmitted on a network or managed by third parties, as in the case of outsourced database management [2].

Integrity is jointly ensured by the access control mechanism and by semantic integrity constraints. Whenever a subject tries to modify some data, the access control mechanism verifies that the subject is authorized to modify the data, and the semantic integrity subsystem verifies that the updated data are correct with respect to a set of *semantic conditions*, referred to as integrity constraints. To protect data from being tampered with while in transit on a network, data can be digitally signed. Finally, the recovery subsystem and the concurrency control mechanism ensure that data are available and correct despite hardware and software failures and accesses from concurrent application programs. Data availability, especially for databases that are available on the Web, can be further strengthened by the use of techniques protecting against denial-of-service attacks.

As the focus of this monograph is on access control models and mechanisms, we do not cover transaction management or semantic integrity. We refer the reader to [40] for an extensive discussion on transaction models, recovery and concurrency control, and to any database textbook for details on semantic integrity. It is important to notice that because the access control mechanism intercepts every access to protected resources, it can also be used to create profiles of accesses by subjects and thus be used in the context of anomaly detection [49] and insider threat protection. Also as current access control systems, like the ones based on XACML [67], are able to take into account a large variety of information including meta-data associated with the data and context information, they can be used for a variety of goals. An example is to grant access to data based on the confidence level of data [30]; in such case, policies specify which is the minimum level of confidence that certain data must have for a given user to access these data for certain tasks. Such policies thus prevent

#### 4 Introduction

the use of incorrect or invalid data for critical tasks. In this example, the metadata used for access control decisions are the confidence levels associated with the data and the goal of the access control policies is not to protect the confidentiality or integrity of the data, but it is to control that users use data that are “good enough” for the tasks they have to perform.

It is also important to note that an access control mechanism must rely for its proper functioning on some authentication mechanism. Such a mechanism identifies users and confirms their identities. Moreover, data may be encrypted when transmitted over networks and when stored on secondary storage. Authentication and encryption techniques are extensively discussed in the current literature on computer network security and we refer the reader to [50] for details on such topics. We will, however, discuss the use of encryption techniques as an approach to implementing access control. We do not attempt to be exhaustive, but try to articulate the rationale for the approaches we believe to be promising.

In the rest of the section, we first present a short historical overview of access control in database systems based on the overview by Bertino and Sandhu [19] (Section 1.1), and then present a road map for the rest of the monograph (Section 1.2).

### 1.1 An Historical Perspective

Early research proposals in the area of access control systems for DBMSs focused on the development of two different classes of models, based on the *discretionary access control (DAC)* policy and on the *mandatory access control (MAC)* policy, respectively. The discretionary access control policy allows subjects to grant authorizations on the data for which they have administration authorization to other subjects. By contrast, the mandatory access control policy regulates accesses to data by subjects on the basis of predefined classifications of subjects and data. Under such a policy even the creator of a data object, like a relation, is not able to grant at its own discretion access authorizations to other subjects. These early access control systems were developed in the framework of relational database systems. The

relational data model, being a declarative high-level model, made it possible to develop declarative languages for the specification of access control policies. The earlier access control models, and the discretionary models in particular, introduced some important principles [36] that set apart access control models for database systems from access control models adopted by operating systems and file systems. The first principle is that access control models for databases should be expressed in terms of the logical data model; thus authorizations for a relational database should be expressed in terms of the logical constructs of the relational data model, that is, relations, relation attributes, and tuples. The second principle is that for databases, in addition to name-based access control, whereby the protected objects are denoted in authorizations by their names, content-based access control has to be supported. Content-based access control allows the system to determine whether to give or deny access to a data item based on the contents of the data item. The development of content-based access control models, which are, in general, based on the specification of conditions against data contents, was made easy in relational databases by the availability of declarative query languages, such as SQL.

In the area of discretionary access control models for relational database systems, the most important early contribution was the development of the System R access control model by Griffith and Wade [35, 41], from which the access control models of current commercial relational DBMSs have been derived. Key features of this model include the concept of decentralized authorization administration, dynamic granting and revocation of authorizations, and the use of views for content-based access control. Also, the initial format of the authorization grant and revoke commands, that are today part of the SQL standard, was developed as part of this model. Subsequent access control models have extended the System R model with a variety of features, such as negative authorization [18], role-based authorization [77], temporal authorization [6], and context-aware authorization [70].

Discretionary access control mechanisms have, however, a major drawback in that they are not able to control how information is propagated and used once it has been accessed by subjects authorized to do so. This weakness makes discretionary access controls vulnerable to

## 6 *Introduction*

malicious attacks, such as Trojan Horses. A Trojan Horse is a program with an apparent or actually useful function, which contains some hidden functions exploiting the legitimate authorizations of the invoking process. Sophisticated Trojan Horses may leak information by means of covert channels, enabling illegal access to data. A covert channel is any component or feature of a system that is misused to encode or represent information for unauthorized transmission, without violating the stated access control policy. A large variety of components or features can be exploited to establish covert channels, including the system clock, operating system interprocess communication primitives, error messages, the existence of particular file names, the concurrency control mechanism, and so forth. The goal of mandatory access control and multilevel database systems was to address such problems through the development of access control models based on data and subject classification, some of which were also incorporated in commercial products. Early mandatory access control models were mainly developed for military applications and were very rigid and suited, at best, for closed and controlled environments. There was considerable discussion in the security community concerning how to eliminate covert channels while maintaining the essential properties of the relational model. The concept of polyinstantiation, that is, the presence of multiple copies with different security levels of a same tuple in a relation, was developed and investigated in this period [79]. Because of the lack of applications and commercial success, companies developing multilevel DBMSs discontinued their production in the early nineties. Covert channels were also widely investigated with considerable focus on the concurrency control mechanisms that, by synchronizing transactions running at different security levels, would introduce an obvious covert channel. However, solutions developed in the research arena to the covert channel problem were not incorporated into commercial products. Interestingly, however, at the beginning of the 2000s, strong security requirements arising in a number of civilian applications have driven a “multilevel security reprise” [80]. Companies have thus reintroduced such systems. The most notable of such systems is Labeled Oracle, a multilevel relational DBMS by Oracle, which has much more flexibility in comparison to earlier multilevel secure DBMSs.



These early approaches to access control have then been extended in the context of advanced DBMSs, such as object-oriented DBMSs and object-relational DBMSs, and other advanced data management systems and applications, such as XML repositories, digital libraries and multimedia data, data warehousing systems, and workflow systems. Most of these systems are characterized by data models that are more expressive than the relational model; typically, these extended models include modeling notions such as inheritance hierarchies, aggregation, and methods. An important requirement for those applications concerns the fact that not only the data need to be protected, but also the database schema may contain sensitive information and, thus, accesses to the schema need to be filtered according to the access control policies. Even though early relational DBMSs did not support access control to the schema information, today several products support such feature. In this respect, access control policies may also need to be protected because they may reveal sensitive information. As such, one may need to define access control policies for objects which are not user data, rather they are other access control policies. Another relevant characteristic of advanced applications is that they often deal with multimedia data, for which the automatic interpretation of contents is much more difficult, and they are, in most cases, accessed by a variety of users external to the system boundaries, such as through Web interfaces. As a consequence both discretionary and mandatory access control models developed for relational DBMSs had to be properly extended to deal with additional modeling concepts. Also, these models often need to rely on metadata information in order to support content-based access control for multimedia data and to support credential-based access control policies to deal with external users. Efforts in this direction include the development of comprehensive access control models for XML [9, 67].

## **1.2 Recent Research Directions**

More recent research directions in the area of access control for database systems have been driven by legal requirements as well as by technology developments. A first research direction is related to privacy-preserving

## 8 Introduction

techniques for databases, an area recently investigated to a considerable extent. Privacy legislation, such as the early Federal Act [26] of 1974, and the more recent Health Insurance Portability and Accountability Act of 1996 (HIPAA) [43] and the Children's Online Privacy Protection Act (COPPA) [25], require organizations to deploy adequate fine-grained access control mechanisms able to control access at the finest granularity possible, that is, at the cell level, and also to take into account additional information, such as the data usage purpose and the data retention period [21]. Privacy is also motivating the development of *oblivious access control*, which is crucial when access control decisions are based by also taking into account (possibly sensitive) information about the subjects seeking accesses to the data. A requirement is thus to be able to enforce access control without disclosing such subject information to the party owning the protected data [22, 81]. A second relevant recent research direction is motivated by the trend of considering databases as a service that can be outsourced to external companies [46]. As outsourced data are encrypted when stored at the service provider, subjects authorized to access the data need to receive the proper keys for decrypting the data. Approaches are thus needed in this context for fine-grained encryption, by which different portions of the data are encrypted with different encryption keys and subjects receive only the keys corresponding to the portions they are entitled to access. A possible approach has been defined in the context of third-party publishing systems for XML data [23]. A third relevant direction is driven by the problem of insider threats, that is, individuals who misuse the data to which they have access to. Protecting from such threats require sophisticated techniques, such as anomaly detection tools able to build profiles of normal data accesses and detect accesses that are anomalous with respect to these profiles. A particular crucial problem in this context is represented by malicious database administrators (DBAs), as a DBA has typically access to the entire database he/she administers. To address this problem solutions have been proposed including the segregation of DBAs from user data, as in the case of the Oracle Database Vault product, and techniques for joint administration of critical database objects.

### 1.3 Organization of the Monograph

We begin with a brief introduction to relevant background notions concerning access control models and mechanisms, and cryptography (Section 2). We then summarize the foundations of access control systems for relational database systems, including the access system developed as part of System R [41] and its extensions (Section 3). As these foundations have been covered in a previous survey by Bertino and Sandhu [19], we keep the presentation very short here and refer the reader to such survey for details. The presentation on the foundations is complemented by some case studies covering access control models and mechanisms supported by current DBMSs (Section 4). In particular, we discuss the Oracle Virtual Private Database mechanism which is an interesting approach to context-based access control and the access control mechanism of SQL Server which has many interesting capabilities, such as the support for roles and negative authorizations. We then cover approaches to fine-grained access control. These approaches allow one to associate access permissions with fine-grained elements within a relation, such as a single tuple or even a single cell

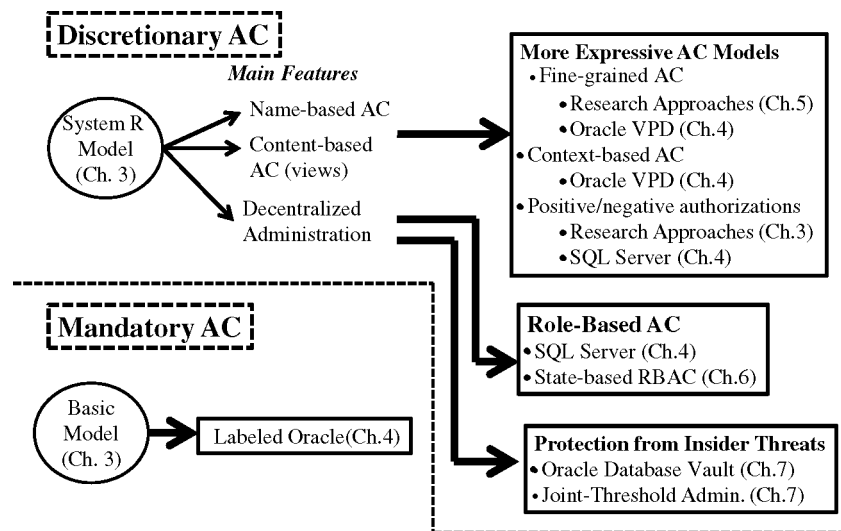


Fig. 1.1 Topics covered in the area of access control for the relational data model.

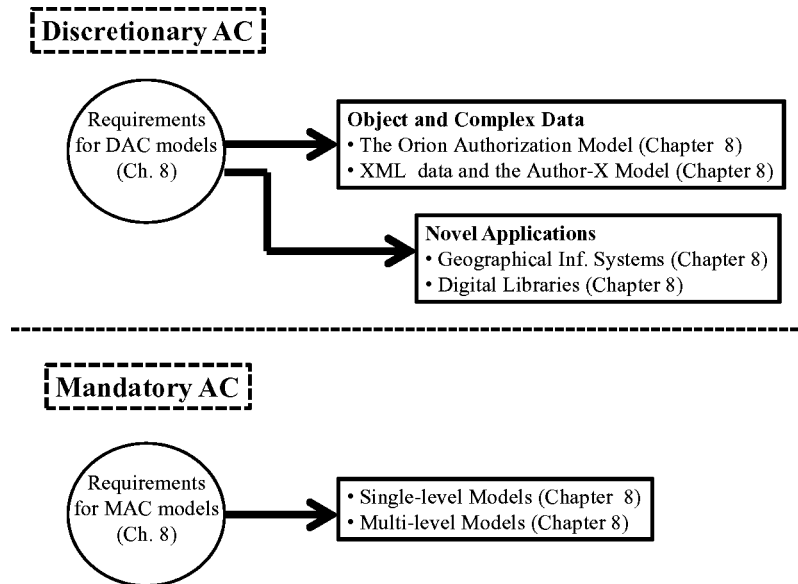


Fig. 1.2 Topics covered in the area of access control for complex data models and selected novel applications.

(Section 5). Fine-grained access control is today a key requirement for information privacy. We then cover more innovative approaches focusing on state-based access control (Section 6), the use of access control mechanisms for protection from insider threats (Section 7), and access control systems for object databases and XML data (Section 8). It is important to remark that approaches and notions developed in the context of object databases, such as those developed for the Orion object-oriented DBMS [74], have been applied to relational DBMSs and also to operating systems. Examples of those approaches and notions include hierarchical authorizations, positive and negative authorizations, and schema protections. We then conclude the paper by discussing the use of cryptography to enforce access control (Section 9), and recent research trends (Section 10). Figures 1.1 and 1.2 provide a high-level description of the relationships among the topics covered in the paper for the relational data model and for more complex data models and selected novel applications, respectively.

## References

---

- [1] R. Agrawal, D. Asonov, M. Kantarcioglu, and Y. Li, "Sovereign joins," in *International Conference on Data Engineering (ICDE)*, 2006.
- [2] R. Agrawal, R. Srikant, and Y. Xu, "Database technologies for electronic commerce," in *Very Large Databases Conference (VLDB)*, 2002.
- [3] ANSI, *Ansi incits 359-2004 for role based access control*. 2004.
- [4] D. Bell and L. LaPadula, "Secure computer systems: Unified exposition and multics interpretation," in *Technical Report*, MTR-2997: Mitre Corporation, 1976.
- [5] A. Belussi, E. Bertino, B. Catania, M. Damiani, and A. Nucita, "An authorization model for geographical maps," in *GIS*, pp. 82–91, New York, NY, USA: ACM, 2004.
- [6] E. Bertino, C. Bettini, E. Ferrari, and P. Samarati, "An access control model supporting periodicity constraints and temporal reasoning," *ACM Transactions on Database Systems (TODS)*, vol. 23, no. 3, pp. 231–285, 1998.
- [7] E. Bertino, P. A. Bonatti, and E. Ferrari, "Trbac: A temporal role-based access control model," *ACM Transactions on Information and System Security (TISSEC)*, vol. 4, no. 3, pp. 191–233, 2001.
- [8] E. Bertino, B. Carminati, E. Ferrari, B. Thuraisingham, and A. Gupta, "Selective and authentic third-party distribution of xml documents," *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 16, no. 10, pp. 1263–1278, 2004.
- [9] E. Bertino, S. Castano, and E. Ferrari, "Securing xml documents with author-x," *IEEE Internet Computing*, vol. 5, no. 3, pp. 21–31, 2001.

148 *References*

- [10] E. Bertino, B. Catania, and E. Ferrari, "A nested transaction model for multilevel secure database management systems," *ACM Transactions on Information and System Security (TISSEC)*, vol. 4, no. 4, pp. 321–370, 2001.
- [11] E. Bertino, B. Catania, E. Ferrari, and P. Perlasca, "A logical framework for reasoning about access control models," *ACM Transaction on Information and System Security (TISSEC)*, vol. 6, pp. 71–127, February 2003.
- [12] E. Bertino and J. Crampton, "Security for distributed systems — foundations of access control," in *Information Assurance: Dependability and Security in Networked Systems*, Morgan Kaufmann, 2008.
- [13] E. Bertino and M. Damiani, "A controlled access to spatial data on web," in *7th AGILE Conference on Geographic Information Science*, pp. 82–91, 2004.
- [14] E. Bertino and E. Ferrari, "Secure and selective dissemination of xml documents," *ACM Transaction on Information Systems Security*, vol. 5, no. 3, pp. 290–331, 2002.
- [15] E. Bertino, E. Ferrari, and A. C. Squicciarini, "Trust-x: A peer-to-peer framework for trust establishment," *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 16, no. 7, pp. 827–842, 2004.
- [16] E. Bertino, S. Jajodia, and P. Samarati, "Database security: Research and practice," *Information Systems*, vol. 20, no. 7, pp. 537–556, 1995.
- [17] E. Bertino, S. Jajodia, and P. Samarati, "A flexible authorization mechanism for relational data management systems," *ACM Transactions on Information Systems*, vol. 17, no. 2, pp. 101–140, 1999.
- [18] E. Bertino, P. Samarati, and S. Jajodia, "An extended authorization model for relational databases," *IEEE Transactions on Knowledge and Data Engineering*, vol. 9, no. 1, pp. 85–101, 1997.
- [19] E. Bertino and R. Sandhu, "Database security — concepts, approaches, and challenges," *IEEE Transactions on Dependable and Secure Computing*, vol. 2, no. 1, pp. 2–19, 1997.
- [20] M. Bykova and M. Atallah, "Succinct specifications of portable document access policies," in *ACM Symposium on Access Control Models and Technologies (SACMAT)*, pp. 41–50, New York, NY, USA: ACM, 2004.
- [21] J.-W. Byun, E. Bertino, and N. Li, "Purpose based access control of complex data for privacy protection," in *ACM Symposium on Access Control Models and Technologies (SACMAT)*, pp. 102–110, 2005.
- [22] J. Camenisch, M. Dubovitskaya, and G. Neven, "Oblivious transfer with access control," in *Proceedings of the 16th ACM Conference on Computer and Communications Security*, pp. 131–140, 2009.
- [23] B. Carminati, E. Ferrari, and E. Bertino, "Securing xml data in third-party distribution systems," in *ACM International Conference on Information and Knowledge Management (CIKM)*, pp. 99–106, 2005.
- [24] S. Chaudhuri, T. Dutta, and S. Sudarshan, "Fine grained authorization through predicated grants," in *Proceedings of the 23rd IEEE International Conference on Data Engineering*, pp. 1174–1183, 2007.
- [25] Children's online privacy protection act of 1998. Available online at <http://www.ftc.gov/ogc/coppa1.htm>. 07 Feb 2009.

- [26] F. T. Commission, Available at [http://www.ftc.gov/foia/privacy\\_act.shtm](http://www.ftc.gov/foia/privacy_act.shtm), Ftc announces settlement with bankrupt website, toysmart.com, regarding alleged privacy policy violations.
- [27] O. Consortium, "Opengis simple features specification for sql," in *Technical Report OGC 99-049*, 1999.
- [28] J. Crampton and G. Loizou, "Administrative scope: A foundation for role-based administrative models," *ACM Transactions on Information and System Security (TISSEC)*, vol. 6, no. 2, pp. 201–231, 2003.
- [29] Create role (transact-sql). Available online at <http://msdn.microsoft.com/en-us/library/ms187936.aspx>. 07 Feb 2010.
- [30] C. Dai, D. Lin, M. Kantarcioglu, E. Bertino, E. Celikel, and B. Thuraisingham, "Query processing techniques for compliance with data confidence policies," in *Secure Data Management Workshop (SDM)*, Springer, 2009.
- [31] M. L. Damiani, E. Bertino, B. Catania, and P. Perlasca, "Geo-rbac: A spatially aware rbac," *ACM Transactions on Information and System Security (TISSEC)*, vol. 10, no. 1, p. 2, 2007.
- [32] Database vault oracle database. Available online at <http://www.oracle.com/database/database-vault.html>. 07 Feb 2010.
- [33] D. E. Denning, "A lattice model of secure information flow," *Communications of the ACM*, vol. 19, no. 5, pp. 236–243, 1976.
- [34] Deny (transact-sql). Available online at <http://msdn.microsoft.com/en-us/library/ms188338.aspx>. 07 Feb 2010.
- [35] R. Fagin, "On an authorization mechanism," *ACM Transactions on Database Systems*, vol. 3, no. 3, pp. 310–319, 1978.
- [36] E. Fernandez, R. Summers, and C. Wood, *Database Security and Integrity*. Addison-Wesley, 1981.
- [37] D. Ferraiolo, D. Kuhn, and R. Chandramouli, *Role-Based Access Control*. Artech House, 2007.
- [38] E. Ferrari, N. R. Adam, V. Atluri, E. Bertino, and U. Capuzzo, "An authorization system for digital libraries," *VLDB Journal*, vol. 11, no. 1, pp. 58–67, 2002.
- [39] R. Gennaro, T. Rabin, S. Jarecki, and H. Krawczyk, "Robust and efficient sharing of rsa functions," *Journal of Cryptology*, vol. 20, no. 3, p. 393, 2007.
- [40] J. Gray and A. Reuter, *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann, 1993.
- [41] P. Griffiths and B. Wade, "An authorization mechanism for a relational database system," *ACM Transactions on Database Systems*, vol. 1, no. 3, pp. 242–255, 1976.
- [42] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman, "Protection in operating systems," *Communications of ACM*, vol. 19, no. 8, pp. 461–471, 1976.
- [43] Health insurance portability and accountability act of 1996. Available online at <http://www.cms.hhs.gov/hipaageninfo/downloads/hipaalaw.pdf>. 07 Feb 2009.
- [44] Incits/iso/iec 9075. sql-99 standard. Available online at <http://webstore.ansi.org/>. 02 Jan 2009.
- [45] Iso 10181-3 access control framework, 1997.

150 *References*

- [46] B. Iyer, S. Mehrotra, E. Mykletun, G. Tsudik, and Y. Wu, "A framework for efficient storage security in rdbms," in *International Conference on Extending Database Technology (EDBT)*, 2004.
- [47] A. Kamra and E. Bertino, "Design and implementation of an intrusion response system for relational databases," *IEEE Transactions on Knowledge and Data Engineering*, vol. 99, no. PrePrints, 2010.
- [48] A. Kamra and E. Bertino, "Privilege states based access control for fine-grained intrusion response," in *Recent Advances in Intrusion Detection (RAID)*, pp. 402–421, 2010.
- [49] A. Kamra, E. Terzi, and E. Bertino, "Detecting anomalous access patterns in relational databases," *VLDB Journal*, vol. 17, no. 5, pp. 1063–1077, 2008.
- [50] C. Kaufman, R. Perlman, and M. Speciner, *Network Security: Private Communication in a Public World*. Prentice-Hall, 2002.
- [51] W. Kim, J. F. Garza, N. Ballou, and D. Woelk, "Architecture of the orion next-generation database system," *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 2, no. 1, pp. 109–124, 1990.
- [52] D. Kincaid and W. Cheney, *Numerical Analysis: Mathematics of Scientific Computing*. Brooks Cole, 2001.
- [53] B. W. Lampson, "Protection," *SIGOPS Operating Systems Review*, vol. 8, no. 1, pp. 18–24, 1974.
- [54] K. LeFevre, R. Agrawal, V. Ercegovic, R. Ramakrishnan, Y. Xu, and D. DeWitt, "Limiting disclosure in hippocratic databases," in *Proceedings of the 30th International Conference on Very Large Data Bases*, pp. 108–119, 2004.
- [55] J. Li and N. Li, "OACerts: Oblivious Attribute Certificates," *IEEE Transactions on Dependable and Secure Computing*, vol. 3, pp. 340–352, 2006.
- [56] N. Li, J.-W. Byun, and E. Bertino, "A critique of the ansi standard on role-based access control," *IEEE Security and Privacy*, vol. 5, no. 6, pp. 41–49, 2007.
- [57] Liberty alliance project, 2001. Available online at <http://www.projectliberty.org>, 07 Feb 2010.
- [58] D. Lin, P. Rao, E. Bertino, and J. Lobo, "An approach to evaluate policy similarity," in *ACM symposium on Access control models and technologies (SACMAT)*, pp. 1–10, New York, NY, USA: ACM, 2007.
- [59] G. Mella, E. Ferrari, E. Bertino, and Y. Koglin, "Controlled and cooperative updates of xml documents in byzantine and failure-prone distributed systems," *ACM Transaction on Information and System Security (TISSEC)*, vol. 9, no. 4, pp. 421–460, 2006.
- [60] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone, *Handbook of Applied Cryptography*. CRC Press, 2001.
- [61] G. Miklau and D. Suciu, "Controlling access to published data using cryptography," in *Proceedings of the 29th international conference on Very large data bases*, pp. 898–909, 2003.
- [62] J. Moss, *Nested Transactions: An Approach to Reliable Distributed Computing*. MIT Press, 1985.
- [63] M. Naor and B. Pinkas, "Oblivious transfer and polynomial evaluation," in *Symposium on Theory of Computer Science (STOC)*, pp. 245–254, 1999.



- [64] R. V. Nehme, E. A. Rundensteiner, and E. Bertino, "A security punctuation framework for enforcing access control on streaming data," in *IEEE International Conference on Data Engineering (ICDE)*, pp. 406–415, Washington, DC, USA: IEEE Computer Society, 2008.
- [65] Q. Ni, E. Bertino, and J. Lobo, "Risk-based access control systems built on fuzzy inferences," in *ACM Symposium on Information, Computer and Communication (ASIACCS)*, 2010.
- [66] Q. Ni, E. Bertino, J. Lobo, and S. B. Calo, "Privacy-aware role-based access control," *IEEE Security and Privacy*, vol. 7, no. 4, pp. 35–43, 2009.
- [67] Oasis consortium, extensible access control markup language (xacml) committee specification, version 1.1. Available online at [http://www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=xacml](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml). 07 Feb 2009.
- [68] Oracle, "Oracle label security administrator's guide 10g release 10g release 2 (10.2) b14267-02," Available online at [http://www.oracle.com/pls/db102/to\\_pdf?pathname=network.10230](http://www.oracle.com/pls/db102/to_pdf?pathname=network.10230) Jan 2010.
- [69] Oracle 11g. Available online at <http://www.oracle.com/index.html>. 07 Feb 2010.
- [70] Oracle database 11g virtual private database. Available online at <http://www.oracle.com/technology/deploy/security/database-security/virtual-private-database/index.html>. 07 Feb 2010.
- [71] Oracle database concepts 11g release 1 (11.1). Available online at [http://download.oracle.com/docs/cd/B28359\\_01/server.111/b28318/datadict.htm](http://download.oracle.com/docs/cd/B28359_01/server.111/b28318/datadict.htm). 03 Jul 2009.
- [72] T. P. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *CRYPTO '91: Proceedings of the 11th Annual International Cryptology Conference on Advances in Cryptology*, pp. 129–140, 1992.
- [73] PostgreSQL global development group. postgresql 8.3 documentation. Available online at <http://www.postgresql.org/docs/8.3/static/sql-grant.html>. 02 Jan 2009.
- [74] F. Rabitti, E. Bertino, W. Kim, and D. Woelk, "A model of authorization for next-generation database systems," *ACM Transactions on Database Systems (TODS)*, vol. 16, no. 1, pp. 88–131, 1991.
- [75] P. Rao, G. Ghinita, E. Bertino, and J. Lobo, "Visualization for access control policy analysis results using multi-level grids," in *IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY)*, pp. 25–28, Washington, DC, USA: IEEE Computer Society, 2009.
- [76] S. Rizvi, A. Mendelzon, S. Sudarshan, and P. Roy, "Extending query rewriting techniques for fine-grained access control," in *Proceedings of the ACM International Conference on Management of Data*, pp. 551–562, 2004.
- [77] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman, "Role-based access control models," *Computer*, vol. 29, no. 2, pp. 38–47, 1996.
- [78] R. Sandhu, D. Ferraiolo, and R. Kuhn, "The nist model for role-based access control: Towards a unified standard," in *ACM Workshop on Role-based Access Control*, pp. 47–63, 2000.
- [79] R. S. Sandhu and F. Chen, "The multilevel relational (mlr) data model," *ACM Transactions on Information and System Security*, vol. 1, no. 1, pp. 93–132, 1998.

152 *References*

- [80] O. S. Saydjari, "Multilevel security: Reprise," *IEEE Security and Privacy*, vol. 2, no. 5, pp. 64–67, 2004.
- [81] N. Shang, M. Nabeel, F. Paci, and E. Bertino, "A privacy-preserving approach to policy-based content dissemination," in *Proceedings of the 26th IEEE International Conference on Data Engineering*, 2010.
- [82] V. Shoup, "Practical threshold signatures," in *EUROCRYPT*, pp. 207–220, 2000.
- [83] Sql server 2008. Available online at <http://www.microsoft.com/sqlserver/2008/en/us/default.aspx>. 07 Feb 2010.
- [84] The postgresql global development group. postgresql 8.3. Available online at <http://www.postgresql.org/>. 07 Jun 2009.
- [85] M. B. Thuraisingham, "Mandatory security in object-oriented database systems," in *Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pp. 203–210, New York, NY, USA: ACM, 1989.
- [86] Q. Wang, T. Yu, N. Li, J. Lobo, E. Bertino, K. Irwin, and J.-W. Byun, "On the correctness criteria of fine-grained access control in relational databases," in *Proceedings of the 33rd International Conference on Very Large Data Bases*, pp. 555–566, 2007.
- [87] J. Widom and S. Ceri, *Active Database Systems: Triggers and Rules For Advanced Database Processing*. Morgan Kaufmann, 1996.
- [88] World wide web consortium, platform for privacy preferences (p3p). Available online at <http://www.w3.org/P3P>. 07 Feb 2010.
- [89] Xml tutorial. Available online at <http://www.w3schools.com/xml/default.asp>. 07 Feb 2010.