

Access Control in Data Management Systems

Synthesis Lectures on Data Management

Editor

M. Tamer Özsu, *University of Waterloo*

Synthesis Lectures on Data Management is edited by Tamer Özsu of the University of Waterloo. The series will publish 50- to 125 page publications on topics pertaining to data management. The scope will largely follow the purview of premier information and computer science conferences, such as ACM SIGMOD, VLDB, ICDE, PODS, ICDT, and ACM KDD. Potential topics include, but not are limited to: query languages, database system architectures, transaction management, data warehousing, XML and databases, data stream systems, wide scale data distribution, multimedia data management, data mining, and related subjects.

Access Control in Data Management Systems

Elena Ferrari

2010

An Introduction to Duplicate Detection

Felix Naumann and Melanie Herschel

2010

Privacy-Preserving Data Publishing: An Overview

Raymond Chi-Wing Wong and Ada Wai-Chee Fu

2010

Keyword Search in Databases

Jeffrey Xu Yu, Lu Qin, and Lijun Chang

2009

© Springer Nature Switzerland AG 2022
Reprint of original edition © Morgan & Claypool 2010

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—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

Access Control in Data Management Systems

Elena Ferrari

ISBN: 978-3-031-00708-8 paperback

ISBN: 978-3-031-01836-7 ebook

DOI 10.1007/978-3-031-01836-7

A Publication in the Springer series

SYNTHESIS LECTURES ON DATA MANAGEMENT

Lecture #4

Series Editor: M. Tamer Özsu, *University of Waterloo*

Series ISSN

Synthesis Lectures on Data Management

Print 2153-5418 Electronic 2153-5426

Access Control in Data Management Systems

Elena Ferrari

University of Insubria, Varese, Italy

SYNTHESIS LECTURES ON DATA MANAGEMENT #4

ABSTRACT

Access control is one of the fundamental services that any Data Management System should provide. Its main goal is to protect data from unauthorized read and write operations. This is particularly crucial in today's open and interconnected world, where each kind of information can be easily made available to a huge user population, and where a damage or misuse of data may have unpredictable consequences that go beyond the boundaries where data reside or have been generated.

This book provides an overview of the various developments in access control for data management systems. Discretionary, mandatory, and role-based access control will be discussed, by surveying the most relevant proposals and analyzing the benefits and drawbacks of each paradigm in view of the requirements of different application domains. Access control mechanisms provided by commercial Data Management Systems are presented and discussed. Finally, the last part of the book is devoted to discussion of some of the most challenging and innovative research trends in the area of access control, such as those related to the Web 2.0 revolution or to the Database as a Service paradigm.

This book is a valuable reference for an heterogeneous audience. It can be used as either an extended survey for people who are interested in access control or as a reference book for senior undergraduate or graduate courses in data security with a special focus on access control. It is also useful for technologists, researchers, managers, and developers who want to know more about access control and related emerging trends.

KEYWORDS

data management systems, data protection, confidentiality, integrity, access control, discretionary access control, mandatory access control, role-based access control, authorization administration

Contents

	Acknowledgments	xi
	Preface	xiii
1	Access Control: Basic Concepts	1
1.1	Introduction	1
1.2	Historical Prospective	1
1.3	Data Protection	2
1.4	Basic Components	3
1.5	Access Control Policies	4
1.6	Access Authorizations	6
1.6.1	Authorization Subjects	6
1.6.2	Authorization Objects	7
1.6.3	Authorization Privileges	7
1.6.4	Authorizations: Further Components	7
1.7	Administration Policies	9
1.8	Access Control Models	10
2	Discretionary Access Control for Relational Data Management Systems	11
2.1	The Access Matrix Model	11
2.2	The System R Access Control Model	13
2.2.1	GRANT Command	14
2.2.2	REVOKE Command	16
2.2.3	Authorization Storage	17
2.2.4	Authorization Management and Enforcement	18
2.2.5	Recursive Revocation	19
2.2.6	Authorizations on Views	22

2.3	DAC Support in SQL	25
2.4	Extensions to the System R Access Control Model	26
2.4.1	Positive and Negative Authorizations	27
2.4.2	Temporal Authorizations	29
2.5	Oracle Virtual Private Database	30
3	Discretionary Access Control for Advanced Data Models	37
3.1	Access Control for Object DMSs	37
3.2	Access Control for XML Data	39
3.2.1	Access Control Requirements	40
3.2.2	Access Control Models	41
3.2.3	Efficiency of Access Control	46
4	Mandatory Access Control	49
4.1	Bell and LaPadula Model	50
4.2	Multilevel Relational Data Model	54
4.3	Mandatory Access Control for Object DMSs	56
4.4	MAC vs DAC	57
4.5	Information-flow Control Models	59
5	Role-based Access Control	61
5.1	The ANSI/INCITS RBAC Standard	62
5.1.1	Core RBAC	62
5.1.2	Hierarchical RBAC	64
5.1.3	Constrained RBAC	65
5.2	RBAC Support in SQL	67
5.3	Role Administration	70
5.4	RBAC Extensions	74
6	Emerging Trends in Access Control	77
6.1	Access Control under the Database as a Service Model	77
6.2	Access Control for Data Stream Management Systems	80

6.3 Access Control in the Web 2.0 Era82

6.3.1 OSN Access Control Requirements82

6.3.2 Proposed Solutions83

6.4 Further Research Directions in Access Control85

Bibliography89

Author’s Biography103

Acknowledgments

A work of this kind depends on the cooperation of many people, students, colleagues, and researchers who have worked with me on topics related to access control. Without their efforts this project would not have been possible. I would like to thank Barbara Carminati who read the final draft of this book and provided valuable comments. Special thanks go to M. Tamer Özsu who gave me the opportunity to write this book and provided many valuable comments to improve its quality. I really appreciate the Morgan & Claypool staff, in particular Executive Editor Diane D. Cerra for help throughout this project and Dr. C.L. Tondo and his crew for handling the production of the camera-ready pages. Last, but not least, special thanks to my family for their great and warm support through this exciting but time-consuming project. In particular, I am really indebted to my little son Tommaso Piero, to which this book is dedicated, with love.

Elena Ferrari
April 2010

Preface

We live in a time of unprecedented opportunities for storing, managing, and analyzing data referring to any kind of information, from personal to business-oriented, recorded by a variety of devices that follow us during our daily activities. This huge amount of information is both a challenge and a risk. Indeed, the availability of this source of information is the basic building block of the idea of the *knowledge society*: a society where knowledge is a major component of any human activity and decisions—big or small—can be taken on the basis of reliable knowledge, distilled from ubiquitous generated data. Moreover, the Web 2.0 revolution and its collaborative tools have made access to data easier by potentially unknown users.

In such a scenario, data become one of the most crucial assets and, as such, their protection from any kind of intrusions, improper modifications, theft, and unauthorized disclosures is a fundamental service. Therefore, there is a strong need of models and mechanisms to protect data managed by any Data Management System (DMS). Due to the open and interconnected digital world we are immersed in today, data protection is much more difficult than in the past, because it is almost impossible to design safe boundaries where data can be confined.

Data security [Ferrari, 2009c] is a broad concept that deals with different aspects of data protection (e.g., authentication, integrity, auditing). Such security properties are usually enforced through a set of *security services*, using a variety of techniques (e.g., encryption, digital signatures, trusted hardware/software).

This book is about one of the key components of the security infrastructure of any data management system, that is, *access control* [Ferrari, 2009a]. Access control aims at preventing unauthorized operations (such as read and write) on the managed data. In this book, we first examine what is needed to control access to data, then we explore the major approaches that have been applied in designing access control mechanisms. As we will see in the reminder of the book, the developments in access control are mainly driven by two factors: the development of new data models and the needs of new applications and environments. The overall goal of these developments is to provide more expressive access control models without compromising the security and efficiency of the system¹.

The book is organized as follows. We start our journey in Chapter 1 by providing the basic concepts on access control that will be developed throughout the book. Then, Chapters 2 and 3 are devoted to discretionary access control. Chapter 2 is about relational data management systems. It revises the most important research proposals as well as discusses the access control support provided by SQL and some of the most innovative features provided by relational DMSs. Chapter 3 is devoted to discretionary access control, however it considers data models beyond the relational one and, in

¹Here, and in what follows, expressivity is related to the set of access control requirements that the access control model possibly supports.

particular, the object and XML data models. We believe that the discussion of access control for XML data is fundamental due to the role played by XML for Web data management. Chapter 4 is devoted to the other big family of access control models, besides discretionary ones, that is, mandatory access control models. Besides describing the most relevant proposals in the field, we discuss the main differences with respect to discretionary access control, the environments that can benefit from mandatory access control, as well as the possible drawbacks that have to be considered. Chapter 5 is devoted to the third major player in the access control area, that is, Role-based Access Control (RBAC). We describe the ANSI/INCITS RBAC standard and some of its recent developments. One of the primary goals of RBAC is to simplify administration of access rights. Therefore, a part of the chapter is devoted to discussing the administration models for RBAC proposed so far. In Chapter 6, we highlight several interesting research issues regarding access control; for instance, how to protect data when they are outsourced to a third party, how to protect data streams that should be securely managed on-the-fly, and how to protect resources and personal information in an On-line Social Network. All these environments pose new and fascinating challenges for what concerns access control that sometimes requires rethinking the way access control has been managed so far.

An heterogeneous audience can benefit from this book. First of all, the book could be used as a reference for senior undergraduate or graduate courses in data security which has a special focus on access control. However, it is also useful for technologists, researchers, managers, and developers who want to know more about access control and emerging trends.

Elena Ferrari
April 2010