

Security of Shared Data in Large Systems: State of the Art and Research Directions

Arnon Rosenthal
The MITRE Corporation
202 Burlington Road /k308
Bedford MA 01730 USA
arnie@mitre.org

Marianne Winslett
University of Illinois
Dept. of Computer Science, 201 N. Goodwin
Urbana, IL 61820 USA
winslett@cs.uiuc.edu

Abstract

The goals of this tutorial are to enlighten the VLDB research community about the state of the art in data security, especially for enterprise or larger systems, and to engage the community's interest in improving the state of the art. The tutorial includes numerous suggested topics for research and development projects in data security.

1. Introduction

Security is increasingly recognized as a key impediment to sharing data in enterprise systems, virtual enterprises, and the semantic web. Yet the topic has not been a focus for mainstream database research, industrial progress in data security has been slow, and (too) much security enforcement is in application code, or else is coarse grained and insensitive to data contents.

The VLDB community is in an excellent position to make significant improvements in the way people think about security policies, due to the community's experience with declarative and logic-based specifications, automated compilation and physical design, and both semantic and efficiency issues for federated systems. These strengths provide a foundation for improving theory and practice.

This tutorial aims to enlighten the VLDB research community about the state of the art in *data* security, especially for enterprise or larger systems, and to engage the community's interest in improving the state of the art. Thus after a very brief look at security basics, the tutorial

focuses on the following questions:

- What is the current state of the art in the real world, with respect to data security?
- What sorts of additional research results are needed to improve the current state of affairs?
- What frameworks can be helpful for architects and researchers tackling these problems?

We will present many open research problems (some graduate student ready, others requiring more formalization) as we move through the sections of the tutorial. For architects, we will suggest unifying concepts and distinctions that they could support, even before the research matures.

Our overall goal for the tutorial is to present material that is not found in any textbook---to show the audience how *they* can help improve the state of the art in data security. Thus the first section of the tutorial, entitled *Basics*, is not intended as a replacement for a security textbook. While a tutorial focused entirely on security basics might be helpful to the SIGMOD community, that is not our goal.

2. About the Presenters

Arnie Rosenthal is a Principal Scientist at MITRE. He has broad interests in problems that arise when data is shared between communities, including a long-term interest in the security issues that arise in data warehouses, federated databases, and enterprise information systems. He has also had a first-hand look at many security problems that arise in large government and military organizations.

Marianne Winslett has been a professor at the University of Illinois since 1987. She started working on database security issues in the early 1990s, focusing on semantic issues in MLS databases. Her interests soon shifted to issues of trust management for data on the web. Trust negotiation is her main current research focus.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

**Proceedings of the 30th VLDB Conference,
Toronto, Canada, 2004**