DATA AND APPLICATIONS SECURITY: PAST, PRESENT AND THE FUTURE

Bhavani Thuraisingham
The National Science Foundation, Arlington, VA and The MITRE Corporation, Bedford, MA

Abstract: This paper first describes the developments in data and applications security with a special emphasis on database security. It also discusses current work including security for data warehouses and e-commerce systems. Then it discuses some of the directions for data and applications security which includes secure semantic web, secure dependable information management, and secure sensor information management. Directions for privacy research are also given.

Key words: Database Security, Secure Semantic Web, Secure Dependable Information Management, Secure Sensor Information Management, Privacy

1. INTRODUCTION

Recent developments in information systems technologies have resulted in computerizing many applications in various business areas. Data has become a critical resource in many organizations, and therefore, efficient access to data, sharing the data, extracting information from the data, and making use of the information has become an urgent need. As a result, there have been many efforts on not only integrating the various data sources scattered across several sites, but extracting information from these databases in the form of patterns and trends has also become important. These data sources may be databases managed by database management systems, or they could be data warehoused in a repository from multiple data sources.

The advent of the World Wide Web (WWW) in the mid 1990s has resulted in even greater demand for managing data, information and knowledge effectively. There is now so much data on the web that managing it with conventional tools is becoming almost impossible. New tools and techniques are needed to effectively manage this data. Therefore, to provide interoperability as well as warehousing between the multiple data sources and systems, and to extract information from the databases and warehouses on the web, various tools are being developed.

As the demand for data and information management increases, there is also a critical need for maintaining the security of the databases, applications and information systems. Data and information have to be protected from unauthorized access as well as from malicious corruption. With the advent of the web it is even more important to protect the data and information as numerous individuals now have access to this data and information. Therefore, we need effective mechanisms for securing data and applications.

This paper will review the developments in data and applications security with a special emphasis on database security. Then it will provide directions for data and applications security. These directions include securing emerging applications such as secure semantic web, secure sensor/stream information processing and privacy.

2. DEVELOPMENTS IN DATABASE SECURITY

Initial developments in database security began in the 1970s. For example, as part of the research on System R at IBM Almaden Research Center, there was a lot of work on access control for relational database systems. About the same time, some early work on multilevel secure database management systems (MLS/DBMS) was reported.

However it was only after the Air Force Summer Study in 1982, that much of the developments on secure database systems began (see [1]). There were the early prototypes based on the integrity lock mechanisms developed at the MITRE Corporation. Later in the mid-1980s pioneering research was carried out at SRI International and Honeywell Inc. on systems such as SeaView and Lock Data Views. Some of the technologies developed by these research efforts were transferred to commercial products by corporations such as Oracle, Sybase, and Informix.

The research in the mid 1980s also resulted in exploring some new areas such as the inference problem, secure object database systems, secure transaction processing and secure distributed/federated database systems. In fact Dr. John Campbell of the National Security Agency stated in 1990 is that one of the important developments in database security was the work of Thuraisingham on the unsolvability of the inference problem [6]. This research then led the way to examine various classes of the inference problem. Throughout the early 1990s, there were many efforts reported on these new types of secure database systems by researchers at organizations such as The MITRE Corporation, Naval Research Laboratory, University of

Milano and George Mason University. In addition much work was also carried out on secure transactions processing.

In the mid 1990s, with the advent of the web, there were many new directions for secure data management research. These included secure workflow systems, secure digital libraries, web security, and secure data warehouses. New technologies such as data mining exacerbated the inference problem as even naive users could now use data mining tools and infer sensitive information. However data mining is also a very important technique for solving many security problems such as intrusion detection and auditing. Therefore the challenge is to carry out data mining but at the same time ensuring that the inference problem is limited. Furthermore, with developments in distributed object systems and e-commerce applications, resulted in developments in secure distributed object systems and secure e-commerce applications. In addition, access control also has received a lot of attention especially in the area of role-based access control (RBAC). A fairly comprehensive survey of data management security is reported in the article by Ferrari and Thuraisingham [11]. In this survey, the authors have given numerous references to research in database security starting around the 1970s until around the late 1990s. Also, many of the publications in secure data management have appeared in the Proceedings of the IFIP 11.3 Working Conferences on Data and Applications Security [15] as well as in the proceedings of various security and database conferences. Numerous researchers all over the world have contributed toward the developments in database and applications security. A listing of all the references is beyond the scope of this paper.

3. DIRECTIONS FOR DATA AND APPLICATIONS SECURITY

3.1 Overview

In section 2 we provided a brief overview of the developments in secure data management over a thirty-year period starting around the early 1970s until the late 1990s. While there have been many developments in web security, there is still a lot of work to be done. Every day we are seeing developments on web data management. For example, standards such as XML (extensible Markup Language), RDF (Resource Description Framework) are emerging. Security for these web standards has to be examined. Also, web services are becoming extremely popular and necessary for many applications and therefore we need to examine secure web services.

Semantic web is a concept that will soon become a reality. We need to examine security issues for the semantic web (see for example, [12]).

Security needs to be examined for new application areas such as bioinformatics, peer-to-peer computing, grid computing, and stream/sensor data management. For example, in the case of bioinformatics, it is important that information about the genes of the individuals be protected from unauthorized users. Peer-to-peer computing has received a lot of attention recently. There are numerous security issues for such systems including secure information sharing and collaboration. Furthermore, data is no longer only in structured databases. Data could be streams emanating from sensors and other sources as well as text, images and video. Security for such data has not received much attention. One also has to make tradeoffs between security, data quality and real-time processing. In other words, we need research on quality of service for information processing. Grid computing is developing rapidly especially to develop infrastructures and tools for scientists and engineers to carry out research. We need to ensure that the grids are secure. Finally we need to ensure that individual privacy is maintained especially when tools for data mining are used to extract information for national security purposes.

In summary, as new technologies emerge there are lot of security issues that need to be examined. Furthermore, as technologies are being developed for organizations and societies such as knowledge management and collaboration tools, we need to examine the security impact on these tools and technologies. While we have made a lot of progress in data and applications security the last three decades, there are many more security issues that need to be examined in the next several decades. To address these issues, there are now various research programs being funded in USA and in Europe. A description of the Cyber Trust Theme funded by the National Science Foundation, which includes data, and applications security can be found in [19]. A discussion of some previous efforts can be found in [26]).

The remaining sections will focus on some of the specific directions in secure semantic web, secure dependable information management, and secure sensor information management and privacy. These were some of the direction discussed during the keynote presentation at the IFIP 11.3 Data and Applications Security Conference in Colorado on August 3, 2003 (see [27]. There are many more areas that need work in security including bioinformatics, geoinformatics, and scientific and engineering informatics in general, peer-to-peer information management, and information management for pervasive and mobile computing. A discussion of the security issues for all of these areas is beyond the scope of this paper. Note also that we have not discussed topics such as critical infrastructure protection and information assurance for databases in this paper. For a detailed discussion of database and applications security, we refer to [34].

3.2 Secure Semantic Web

While the current web technologies facilitate the integration of information from a syntactic point of view, there is still a lot to be done to integrate the semantics of various systems and applications. That is, current web technologies depend a lot on the human-in-the-loop for information integration. Tim Berners Lee, the father of WWW, realized the inadequacies of current web technologies and subsequently strived to make the web more intelligent. His goal was to have a web that will essentially alleviate humans from the burden of having to integrate disparate information sources as well as to carry out extensive searches. He then came to the conclusion that one needs machine understandable web pages and the use of ontologies for information integration. This resulted in the notion of the semantic web [17].

Tim Berners Lee has specified various layers for the semantic web. At the lowest level one has the protocols for communication including TCP/IP (Transmission Control Protocol/Internet protocol), HTTP (Hypertext Transfer Protocol) and SSL (Secure Socket Layer). The next level is the XML (eXtensible Markup Language) layer that also includes XML schemas. The next level is the RDF (Resource Description Framework) layer. Next comes the Ontologies and Interoperability layer. Finally at the highest-level one has the Trust Management layer. For the semantic web to operate successfully we need to ensure that security is maintained within and across all of the layers. In this section we provide an overview of the security issues for the semantic web.

As stated earlier, logic, proof and trust are at the highest layers of the semantic web. That is, how can we trust the information that the web gives us? Closely related to trust is security. However security cannot be considered in isolation. That is, there is no one layer that should focus on security. Security cuts across all layers and this is a challenge. For example, consider the lowest layer. One needs secure TCP/IP, secure sockets, and secure HTTP. There are now security protocols for these various lower layer protocols. One needs end-to-end security. That is, one cannot just have secure TCP/IP built on untrusted communication layers. That is, we need network security. Next layer is XML and XML schemas. One needs secure XML (see [4], [5]). That is, access must be controlled to various portions of the document for reading, browsing and modifications. There is research on securing XML and XML schemas. The next step is securing RDF. Now with RDF not only do we need secure XML, we also need security for the interpretations and semantics. For example under certain context, portions of the document may be Unclassified while under certain other context the document may be Classified. As an example one could declassify an RDF document once the war is over. Lot of work has been carried out security constraint processing for relational databases. One needs to determine whether these results could be applied for the semantic web (see [23]).

Once XML and RDF have been secured the next step is to examine security for ontologies and interoperation. That is, ontologies may have security levels attached to them. Certain parts of the ontologies could be Secret while certain other parts may be Unclassified. The challenge is how does one use these ontologies for secure information integration? Researchers have done some work on the secure interoperability of databases. We need to revisit this research and then determine what else needs to be done so that the information on the web can be managed, integrated and exchanged securely.

Closely related to security is privacy. That is, certain portions of the document may be private while certain other portions may be public or semi-private. Privacy has received a lot of attention recently partly due to national security concerns. Privacy for the semantic web may be a critical issue, That is, how does one take advantage of the semantic web and still maintain privacy and sometimes anonymity.

We also need to examine the inference problem for the semantic web. Inference is the process of posing queries and deducing new information. It becomes a problem when the deduced information is something the user is unauthorized to know. With the semantic web, and especially with data mining tools, one can make all kinds of inferences. That is the semantic web exacerbates the inference problem (see [24]). Recently there has been some research on controlling unauthorized inferences on the semantic web. We need to continue with such research (see for example, [10]).

Security should not be an afterthought. We have often heard that one needs to insert security into the system right from the beginning. Similarly security cannot be an afterthought for the semantic web. However, we cannot also make the system inefficient if we must guarantee one hundred percent security at all times. What is needed is a flexible security policy. During some situations we may need one hundred percent security while during some other situations say thirty percent security (whatever that means) may be sufficient.

3.3 Secure Dependable Information Management

Dependable information systems are systems that are secure, survivable, fault tolerant and can process data in real-time. There is quality of service tradeoffs that one needs to make to build dependable information systems. Various information systems including the semantic web have to be dependable. That is, these systems need to be secure, survivable and process data in real-time if needed.

Our previous experience on building next generation command and control systems using real-time objects will form the basis for the discussion on dependable information management. Our infrastructure for the dependable system will consist of objects that have to perform many functions such as interprocess communication as well as memory management. These objects have to incorporate security, real-time processing and fault tolerant computing capabilities. The data managers have to process queries, execute transactions and manage data that may be transient. Data may also be mined to extract information.

The data/information manager will manage various types of data including track data and data for multi-sensor information integration and fusion. The data/information manager will be hosted as an application on the object-based infrastructure. Processing may be distributed among multiple nodes and there has to be coordination between the nodes. Essentially we are proposing a distributed information management system. One could also think of this system as a peer-to-peer system where the nodes are peers that have to work together to solve a problem.

For many applications timing constraints may be associated with data processing. That is, the data may have to be updated within a certain time or it may be invalid. There are tradeoffs between security and real-time processing. That is, it takes time to process the access control rules and as a result, the system may miss the deadlines. Therefore, we need flexible security policies. In certain cases real-time processing may be critical. For example if we are to detect anomalies from the time a person checks in at the ticket counter until he boards the airplane, then this anomaly has to be detected within a certain time. In this case we may have to sacrifice some degree of security. In other cases, we may have a lot of time to say analyze the data and in this case only authorized individuals may access the data.

One possibility for developing dependable infrastructures and data managers for sensor networks is to follow the approach we have developed for real-time command and control systems such as AWACS (Airborne Warning and Control System). Here we developed an infrastructure consisting of a real-time object request broker and services using commercial real-time operating systems. We then developed a real-time data manager and applications hosted on the infrastructure. We used object technology for integrating the various components. We also showed how such an infrastructure could be used to migrate legacy applications (see [3], [25]).

We can take a similar approach to build an integrated system for dependable information systems. We need appropriate operating systems and infrastructures possibly based on object request brokers. We need to host dependable data managers and applications such as multi-sensor data integration and fusion on the infrastructures. We need to ensure that the system is secure and survivable. We also need to ensure that the infrastructure is secure. That is, security has to be built into the system and not considered as an after-thought. Essentially we are proposing a layered architecture for dependable information management. The infrastructure consists of middleware possibly based on object request brokers for sensors. The objects that constitute the middleware include objects for interprocess communication, memory management, and support for data fusion and aggregation. On top of the infrastructure we host a dependable data manager. We also need to examine both centralized and distributed architectures for dependable data management. On the one hand we can aggregate the data and send it to a centralized data management system or we can develop a full-fledged distributed data management system. We need to conduct simulation studies and determine tradeoffs between various architectures and infrastructures.

As part of our work on evolvable real-time command and control systems we examined real-time processing for object request brokers and we were instrumental in establishing a special interest group, which later became a task force within the Object Management Group (OMG). Subsequently commercial real-time object request brokers were developed. The question is, do we need special purpose object request brokers for dependable networks? Our challenge now is to develop object request brokers for dependable systems. We need to examine special features for object request brokers for managing data for dependable systems.

Another aspect that is important is end-to-end dependability. This includes security, real-time processing and fault tolerance for not only all of the components such as infrastructures, data managers, networks, applications and operating systems; we also need to ensure the dependability of the composition of the entire system. This will be a major challenge even if we consider security, real-time processing and fault tolerant computing individually. Integrating all of them will be a complex task and will have to address many research issues and challenges.

3.4 Secure Sensor Information Management

We need to conduct research on security for sensor databases and sensor information systems. Note that sensor information management is one aspect of dependable information management. For example, can we apply various access control techniques for sensor and stream databases? That is, can we give access to the data depending on the roles of the users such as the air port security officer has access to all of the sensor data emanating from the sensors while the airport ticketing agent may have limited access to certain sensor data. Another challenge is granting access to aggregated data. Individual data may be unclassified while the aggregated data may be highly sensitive. This is in a way a form of the inference problem in database systems. Note that inference is the process of posing queries and obtaining unauthorized information from the legitimate responses received. Due to the aggregation and fusion of sensor data, the security levels of the aggregated data may be higher than those of the individual data sets. We also need to be aware of the privacy of the individuals. Much of the sensor data may be about individuals such as video streams about activities and other personal information. This data has to be protected from the general public and from those who are unauthorized to access the data. We have looked at privacy as a subset of security (see for example, [28]). There is also research on privacy preserving data mining and the techniques have to be examined for sensor data mining.

Finally we need to examine security policies for sensor data. These security policies may be dynamic and therefore we need to develop ways to enforce security constraints that vary with time. We also need techniques for integrating security policies especially in a networked and distributed environment. For example, different policies may apply for different sensor databases. These policies have to be integrated when managing distributed databases. One of the major questions here is what are the special considerations for security for sensor and stream data? Do the access control models that have been developed for business data processing applications work for stream data? We need to start a research program on secure sensor networks and secure sensor information management. Some preliminary directions are given in [28].

We need end-to-end security. That is, the infrastructures, data managers, applications, networks and operating systems for sensor information management have to be secure. The Object Management Group has developed standards for securing object request brokers. We need to take advantage of such developments. However, we need to examine security issues specific to object request brokers for sensor information management. We also need to examine composability of the various secure components. For example, are the interfaces between the various components satisfying the security properties? How can we integrate or compose the security policies of the various components. There is little work reported on securing large-scale information systems. Now, we not only have to examine security for such systems, we also need to examine security for such systems that manage and process sensor data. In addition, we need not only secure sensor networks but also sensor networks that are dependable and survivable.

One approach is to develop various features incrementally for data managers and middleware. However this often means that at the end some features are left out. For example, if we build components and examine only security and later on examine real-time processing, then it may be difficult to incorporate real-time processing into the policy. This means that we need to examine all of the features simultaneously. This means security engineering has to be integrated with software engineering.

Other issues include fault tolerant sensors and survivable sensors. Much work has been carried out on fault tolerant data management. We need to examine the fault tolerant data processing techniques for sensor data. Furthermore, these sensor databases have to survive from failures as well as from malicious attacks. Many of our critical infrastructures such as our telephones, power lines, and other systems have embedded sensors. The data emanating from these sensors may be corrupted maliciously or otherwise. For example, how can we ensure that the aggregate data is valid even if the components that constitute the aggregate data may be corrupted? Some directions are given in [20].

3.5 Privacy in Databases

With the World Wide Web, there is now an abundance of information about individuals that one can obtain within seconds. This information could be obtained through mining or just from information retrieval. Data mining is the process of posing queries and extracting information previously unknown machine learning and other reasoning techniques (see [24]). Now, data mining is an important technology for many applications. However data mining also causes privacy concerns as users can now put pieces of information together and extract information that is sensitive or private. Therefore, one needs to enforce controls on databases and data mining tools. That is, while data mining is an important tool for many applications, we do not want the information extracted to be used in an incorrect manner. For example, based on information about a person, an insurance company could deny insurance or a loan agency could deny loans. In many cases these denials may not be legitimate. Therefore, information providers have to be very careful in what they release. Also, data mining researchers have to ensure that privacy aspects are addressed.

We are beginning to realize that many of the techniques that were developed for the past two decades or so on the inference problem can now be used to handle privacy. One of the challenges to securing databases is the inference problem (see [1]). Inference is the process of users posing queries and deducing unauthorized information from the legitimate responses that they receive. This problem has been discussed quite a lot over the past two decades (see [21], [18], [14]). However, data mining makes this problem worse. Users now have sophisticated tools that they can use to get data and deduce patterns that could be sensitive. Without these data mining tools, users would have to be fairly sophisticated in their reasoning to be able to deduce information from posing queries to the databases. That is, data mining tools make the inference problem quite dangerous [7]. While the inference problem mainly deals with secrecy and confidentiality we are beginning to see many parallels between the inference problem and what we now call the privacy problem.

When Inference problem is considered to be a privacy problem, then we can use the inference controller approach to address privacy. For example, we can develop privacy controllers similar to our approach to developing inference controllers [22]. Furthermore, we can also have different degrees of privacy. For example, names and ages together could be less private while names and salaries together could be more private. Names and healthcare records together could be most private. One can then assign some probability or fuzzy value associated with the privacy of an attribute or a collection of attributes. Lot of work has been carried out on the inference problem in the past. We have conducted some research on applying the techniques for handling the inference problem and apply them for the privacy problem. We discuss some of the directions here.

In one area of research we have designed a privacy enhanced database management system (PE-DBMS) where users with different roles access and share a database consisting of data at different privacy levels. A powerful and dynamic approach to assigning privacy levels to data is one, which utilizes privacy constraints. Privacy constraints provide a privacy policy. They can be used to assign privacy levels to the data depending on their content and the context in which the data is displayed. They can also be used to dynamically reclassify the data. In other words, the privacy constraints are useful for describing privacy-enhanced applications. In [29] we have described the design if a PE-DBMS that processes privacy constraints during database design as well as during query and update operations. Recently Jajodia et al have proposed an approach for release control in databases (see [16]. The idea here is to process constraints after the response is computed but before the document is released. This approach could also be adapted for privacy. Another direction for privacy in databases includes using semantic data models for describing the application, reasoning about the application and detecting privacy violations. Some preliminary results are reported in [30].

We have also proved that the general privacy problem is unsolvable and the next step is to examine the complexity of the privacy problem (see [31]). Essentially we have examined our previous work on the inference problem and adapted it to handle the privacy problem.

There is now research at various laboratories on privacy enhanced/sensitive data mining (e.g., Agrawal at IBM Almaden, Gehrke at Cornell University and Clifton at Purdue University, see for example [2], [8], [13], [9]). The idea here is to continue with mining but at the same time ensure privacy as much as possible. For example, Clifton has proposed the use of the multiparty security policy approach for carrying out privacy sensitive data mining. While there is some progress we still have a long way to go. A summary of the developments in privacy preserving data mining is given in [32]. For a detailed discussion of data mining for national security and the implications for privacy we refer to [33].

3.6 Other Security Considerations

As we have mentioned in section 3.1, numerous technologies have emerged over the past few years. These include tools for electronic commerce, gene extraction, geospatial information management, managing moving objects, grid computing, peer-to-peer information management, managing virtual societies and organizations, and pervasive and context aware computing. Security as well as privacy must be examined for these tools and technologies for them to be useful and effective. For example, we could have the best grid computing tools, but if these tools cannot maintain security and privacy at an acceptable level, then they may not be of much use.

We also need to conduct research on usability vs. security. For example, security enforcement techniques may be quite complicated that it takes time to enforcement them. As a result, one could miss deadlines and/or lose valuable information. In other words, often there is a trade-off between security and technologies. How do we evaluate the economic aspects for incorporating security. That is, would it be economically infeasible to enforce all of the access control policies? Do we have to live with partial security? On the other hand by not having security, banks and financial organizations could lose millions of dollars through fraud. Here cyber security helps economic security. In other words, while security may be expensive to incorporate, for many applications we could lose millions of dollars by not enforcing security. We need to carry out trade off studies and examine the return on investment for incorporating security.

4. SUMMARY

This paper has briefly discussed the developments in data and applications security and then described some directions. We first provided a brief overview of the some of the developments in multilevel secure database management systems including relational systems, objects systems, transaction processing, and the inference problem. In the area of directions, we listed many topics such as secure grid computing, secure bioinformatics and secure geospatial information systems and then focussed on some selected topics such as secure semantic web, secure dependable computing, secure sensor information management and privacy in databases. We chose these topics as they are relevant to our current research and also they elaborate on the some of the issues covered during the keynote presentation.

As we have stressed in this paper, all of the topics listed here are important. That is, we need to conduct research on secure pervasive computing as well as examine the economic aspects of incorporating security. We also need to examine the security impact on knowledge management tools for societies and organizations. While we have made good progress, there is still a lot of research that needs to be done on data and applications security. Organizations such as the National Science Foundation now have focussed programs on Cyber security, which includes data and applications security. We can expect some significant developments on security for new generation information management technologies over the next several years.

5. Disclaimer

The views and conclusions expressed in this paper are those of the author and do not reflect the policies of the National Science Foundation or of the MITRE Corporation.

6. Acknowledgements

I thank the National Science Foundation and the MITRE Corporation for their support of my research on data and applications security.

7. REFERENCES

[1] Air Force Summer Study Report on Multilevel Secure Database Systems, Washington DC, 1983.

[2] Agrawal, R, and R. Srikant, “Privacy-preserving Data Mining,” Proceedings of the ACM SIGMOD Conference, Dallas, TX, May 2000.

[3] Bensley, E., Thuraisingham, B., et al., Design and Implementation of an Infrastructure and Data Management for Evolvable Real-time Command and Control Systems, proceedings IEEE WORDS, Laguna beach, CA, February 1996.

[4] Bertino, E., E. Ferrari, et al, Access Control for XML Documents, Data and Knowledge Engineering, 2002.

[5] Bertino, E., B. Carminati, E. Ferrari, et al, Secure Third Party Publication of XML Documents, Accepted for publication in IEEE Transactions on Knowledge and Data Engineering, 2003.

[6] Proceedings of the National Computer Security Conference, Developments in Database Security, by Dr. John Campbell, 1990.

[7] Clifton, C. and D. Marks, “Security and Privacy Implications of Data Mining”, Proceedings of the ACM SIGMOD Conference Workshop on Research Issues in Data Mining and Knowledge Discovery, Montreal, June 1996.

[8] Clifton, C., M. Kantarcioglu and J. Vaidya, “Defining Privacy for Data Mining,” Purdue University, 2002 (see also Next Generation Data Mining Workshop, Baltimore, MD, November 2002).

[9] Evfimievski, A., R. Srikant, R. Agrawal, and J. Gehrke, “Privacy Preserving Mining of Association Rules,” In Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Edmonton, Alberta, Canada, July 2002.

[10] Farkas, C., Inference Problem for the Semantic Web, Proceedings of the IFIP Conference on Data and Applications Security, Colorado, August 2003.

[11] Ferrari E., and B. Thuraisingham, “Secure Database Systems,” in Advances in Database Management by Artech House, 2000 (Editors: M. Piattini, O. Diaz).

[12] Ferrari, E., and B. Thuraisingham, “Security and Privacy for Web Databases and Web Services,” Proceedings of the EDBT Conference, Crete, March 2004.

[13] Gehrke, J., “Research Problems in Data Stream Processing and Privacy-Preserving Data Mining,” Proceedings of the Next Generation Data Mining Workshop, Baltimore, MD, November 2002.

[14] Hinke T., "Inference and Aggregation Detection in Database Management Systems," Proceedings of the Security and Privacy Conference, Oakland, CA, April 1988.

[15] Proceedings of the IFIP 11.3 Working Conference on Database Security, 1987 – 2003 (North Holland, Chapman and Hall, Kluwer).

[16] Jajodia, S., et al, "Release Control in Documents," Proceedings of the IFIP Database Integrity and Control Conference, Lausanne, Switzerland, November 2003.

[17] Berners Lee, T., et al., "The Semantic Web," Scientific American, May 2001.

[18] Morgenstern, M., "Security and Inference in Multilevel Database and Knowledge Base Systems," Proceedings of the ACM SIGMOD Conference, San Francisco, CA, June 1987.

[19] Cyber Trust Theme, National Science Foundation, http://www.nsf.gov/pubsys/ods/getpub.cfm?ods_key=nsf04524

[20] A. Perrig, "SIA: Secure Information Aggregation in Sensor Networks," CMU Report, 2003.

[21] Thuraisingham, B., "Multilevel Security for Relational Database Systems Augmented by an Inference Engine," Computers and Security," December 1987.

[22] Thuraisingham, B., W. Ford and M. Collins, "Design and Implementation of a Database Inference Controller," Data and Knowledge Engineering Journal, December 1993.

[23] Thuraisingham B. and W. Ford, "Security Constraint Processing in a Distributed Database Management System," IEEE Transactions on Knowledge and Data Engineering, 1995.

[24] Thuraisingham, B., "Data Mining: Technologies, Techniques, Tools and Trends," CRC Press, 1998.

[25] Thuraisingham, B. and J. Maurer, "Information Survivability for Real-time Command sand Control Systems," IEEE Transactions on Knowledge and Data Engineering, January 1999.

[26] Thuraisingham, B., "Data and Applications Security: Developments and Directions," Proceedings of the IEEE COMPSAC Conference, Oxford, UK, August 2002.

[27] Thuraisingham, B., "Data and Applications Security: Developments and Directions," Keynote Presentation, IFIP 11.3 Conference on Data and Applications Security, Estes Park, Colorado, August 3, 2003.

[28] Thuraisingham, B., "Security and Privacy for Sensor Databases," Accepted for publication in Sensor Letters, 2003.

[29] Thuraisingham, B., “Privacy Constraint Processing in a Privacy Enhanced Database System,” Accepted for publication in Data and Knowledge Engineering Journal, 2003.

[30] Thuraisingham, B., “Semantic Data Modeling for Privacy Control in Databases,” Submitted for Publication, 2003.

[31] Thuraisingham, B., “On the Unsolvability of the Privacy Problem in Databases,” Submitted for Publication, 2003.

[32] Thuraisingham, B., “Privacy-Preserving Data Mining: Developments and Directions” Accepted for publication in Journal of Database Management, Journal, 2003.

[33] Thuraisingham, B., “Web Data Mining and Applications in Business Intelligence and Counter-terrorism, CRC Press,” June 2003.

[34] Thuraisingham, B., “Database and Applications Security: Integrating Databases and Information Security,” CRC Press, to appear, 2004.