Secure Knowledge Management: Confidentiality, Trust, and Privacy

Elisa Bertino, *Fellow, IEEE*, Latifur R. Khan, Ravi Sandhu, *Fellow, IEEE*, and Bhavani Thuraisingham, *Fellow, IEEE*

***Abstract*—Knowledge management enhances the value of a corporation by identifying the assets and expertise as well as efficiently managing the resources. Security for knowledge management is critical as organizations have to protect their intellectual assets. Therefore, only authorized individuals must be permitted to execute various operations and functions in an organization. In this paper, secure knowledge management will be discussed, focusing on confidentiality, trust, and privacy. In particular, certain access-control techniques will be investigated, and trust management as well as privacy control for knowledge management will be explored.**

***Index Terms*—Data mining, privacy, role-based access control (RBAC), secure knowledge management, security policy, semantic web, trust negotiation (TN), usage control (UCON).**

I. Introduction

KNOWLEDGE management is about corporations sharing their resources and expertise, as well as building intellectual capital so that they can increase their competitiveness. While knowledge-management practices have been around for decades, it is only with advent of the web that knowledge management has emerged as a technology area. Corporations with Intranets promote knowledge management so that the employees can learn about various advances in technology, get corporate information, and find the expertise in the corporation. Furthermore, when experts leave the corporation through retirement or otherwise, it is important to capture their knowledge and practices so that the corporation does not lose the valuable information acquired through many years of hard work [15].

One of the challenges in knowledge management is maintaining security. Knowledge management includes many technologies such as data mining, multimedia, collaboration, and the web. Therefore, security in web data management, multimedia systems, and collaboration systems all contribute toward securing knowledge-management practices. In addition, one needs to protect the corporation's assets such as its intellectual property. Trade secrets have to be kept highly confidential so that competitors do not have any access to it. This means one needs to enforce some form of access control such as role-based access control (RBAC), credential mechanism, and encryption.

To have secure knowledge management, we need to have secure strategies, processes, and metrics. Metrics must include support for security-related information. Processes must include secure operations. Strategies must include security strategies. The creator may specify to whom the knowledge can be transferred when knowledge is created. The manager of the knowledge may enforce additional access-control techniques. Knowledge sharing and knowledge transfer operations must also enforce access control and security policies. Secure knowledge-management architecture may be built around the corporation's Intranet.

While this paper provides an overview of secure knowledge management, it focuses mainly on confidentiality, trust, and privacy aspects. Section II provides an overview of knowledge management as well as aspects of secure knowledge management. Applying various access-control policies such as RBAC and usage control (UCON) is discussed in Section III. Aspects of trust management and negotiation are discussed in Section IV. Privacy issues are discussed in Section V. The paper is concluded in Section VI.

It should be noted that there are several aspects to secure knowledge management. We discuss confidentiality, trust, and privacy management. Furthermore, for each aspect, several techniques have been proposed. For example, various access-control policies such as read/write policies have been designed and implemented. We will focus on RBAC and UCON as they are emerging as two of the prominent access-control techniques and are natural for use in a corporate environment. RBAC is now being enforced in many commercial products and is a widely used standard. Models such as UCON encompass RBAC models and are becoming popular. Trust management and negotiation is inherent to knowledge management. For example, in an organization, a vice president may be authorized to receive the information, but the president may not have sufficient trust in the vice president to share the sensitive information. Finally, privacy is becoming critical for organizational knowledge management and data sharing. Furthermore, important knowledge-management technologies such as data mining and the semantic web have inferencing capabilities, and as a result, privacy as well as confidentiality may be compromised.

Researchers are working on secure knowledge management that complements the ideas we have presented in this paper. An excellent introduction to secure knowledge management is given in [25]. Work on trust management and policy-driven approaches for the semantic web have been reported in [7], [9], [11], and [12].

Manuscript received February 1, 2005; revised July 1, 2005. This paper was recommended by Guest Editors H. R. Rao and S. J. Upadhyaya.

E. Bertino is with the Department of Computer Sciences, Purdue University, West Lafayette, IN 47907 USA (e-mail: bertino@cs.purdue.edu).

L. R. Khan and B. Thuraisingham are with The University of Texas, Dallas, Richardson, TX 75083 USA (e-mail: bhavani.thuraisingham@utdallas.edu).

R. Sandhu is with the Industrial and Systems Engineering (ISE) Department, George Mason University, Fairfax, VA 22030 USA.

Digital Object Identifier 10.1109/TSMCA.2006.871796

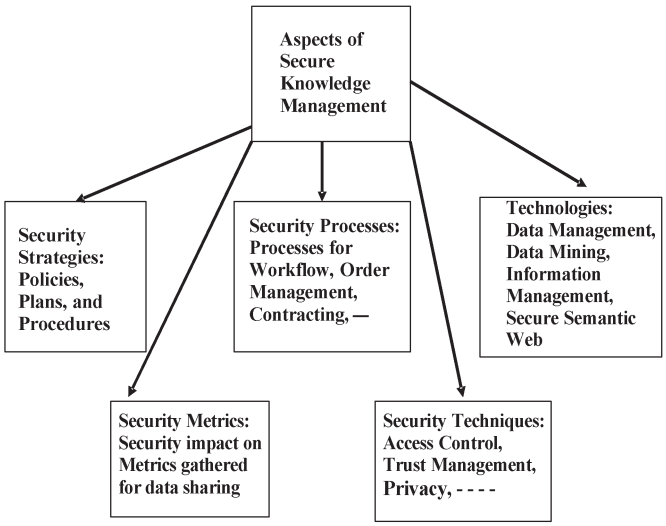

Fig. 1. Aspects of secure knowledge management.

II. Secure Knowledge Management

As stated in Section I, secure knowledge management consists of secure strategies, processes, and metrics. In addition, security technologies such as the secure semantic web and privacy-preserving data mining are technologies for secure knowledge management. Security techniques for knowledge management include access control and trust management. Fig. 1 illustrates the various aspects of secure knowledge management. We will describe each component in this section.

Security strategies for knowledge management include the policies and procedures that an organization sets in place for secure data and information sharing as well as protecting the intellectual property. Some of the questions that need to be answered include how often should knowledge be collected? How often should the organization conduct audit strategies? What are the protection measures that need to be enforced for secure knowledge sharing? Secure knowledge-management strategies should be tightly integrated with business strategies. That is, if by enforcing intellectual-property protection the organization is going to be unprofitable, then the organization has to rethink its secure knowledge-management strategy.

Secure processes for knowledge management include secure workflow processes as well as secure processes for contracting, purchasing, and order management. Security has to be incorporated into the business processes for workflow, contracting, and purchasing. For example, only users with certain credentials can carry out various knowledge-management processes. Metrics for secure knowledge management should focus on the impact of security on knowledge-management metrics. Some examples of knowledge-management metrics include the number of documents published, number of conferences attended, or the number of patents obtained. When security is incorporated, then the number of documents published may decrease as some of the documents may be classified. Organizations should carry out experiments determining the impact of security on the metrics gathered.

Security techniques include access control, UCON, trust management, as well as privacy control. These techniques are enforced at all stages of knowledge-management processes. Secure knowledge-management technologies include data mining, the semantic web, as well as technologies for data and information management. The component technologies have to be secure if we are to ensure secure knowledge management.

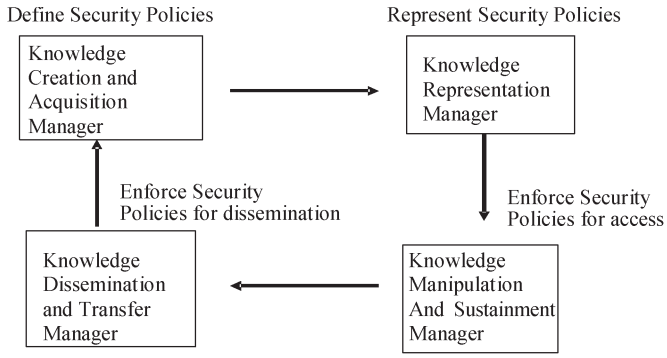

Fig. 2. Secure knowledge-management architecture.

Fig. 2 illustrates an architecture for secure knowledge management. The components of the architecture are a secure knowledge-creation manager, secure knowledge-representation manager, a secure knowledge manipulation and sustainment manager, and a secure knowledge dissemination and transfer manager. The secure knowledge-creation task includes creating knowledge as well as specifying security policies enforced based on the knowledge. Secure knowledge-representation tasks include representing the knowledge as well as policies in a machine-understandable format. Knowledge-representation languages such as rules and frames as well as some of the more recent semantic-web languages such as resource descriptive framework (RDF) and ontology languages are appropriate for knowledge and policy representation. Secure knowledge-manipulation tasks include querying and updating the knowledge base. In addition, the knowledge gained has to be sustained as long as possible. Various processes have to be in place to sustain the knowledge securely. Finally, secure knowledge dissemination and transfer task includes disseminating and transferring the knowledge to authorized individuals.

This section has provided an overview of the various aspects of secure knowledge management. The remainder of the paper will discuss certain security techniques for confidentiality, trust, and privacy.

III. Optional Information for Computer Geeks on Creation of Electronic Image Files

As we have stated, various access-control techniques may be applied for knowledge management to ensure confidentiality. These could be simple read–write policies, role-based policies, and UCON policies. We have chosen to discuss RBAC and UCON as we believe that they are important for secure knowledge management. RBAC is a standard implemented in several commercial products, and UCON encompasses RBAC.

1) RBAC for Knowledge Management: The concept of RBAC has emerged in the past decade as a widely deployed and highly successful alternative to conventional discretionary and mandatory access controls. The principal idea in RBAC is

that users and permissions are assigned to roles. Users acquire permissions indirectly via roles. For example, Alice is assigned to the analyst role and gets the permissions of the analyst role. This remarkably simple idea has many benefits and elaborations. Administration of authorization is much simplified in RBAC. Separation of user–role assignment and permission–role assignment facilitates different business processes and administrators for these tasks. Modifications to the permissions of a role immediately apply to all users in the role. Users are easily deassigned and assigned roles as their job functions and responsibilities change. There are two major elaborations of the simple RBAC concept. One elaboration is to have hierarchical roles, such as senior analyst and junior analyst. The senior analyst automatically inherits the permissions assigned to the junior analyst. This further simplifies administration. The second elaboration is to have separation of duty and other constraints. For example, we may require the roles of analyst and mission specialist to be mutually exclusive, so the same user cannot be assigned to both roles.

Due to strong commercial interest by vendors and users in RBAC over time, the RBAC model evolved into a National Institute of Standards and Technology/American National Standards Institute (NIST/ANSI) standard model first introduced in 2001 [8] and formally adopted as an ANSI standard in 2004. Even though there is agreement on a core standard for RBAC, much work remains to be done to truly exploit the potential of this technology for various applications including knowledge management. RBAC is especially relevant to the protection of information in a local environment as well as in a global environment across a coalition. The NIST/ANSI standard model only captures aspects of RBAC that were mature for standardization. It explicitly lists a number of items on which the standard is silent. These include administration of roles and cross-organizational roles. These open issues are central to deployment of RBAC in complex environments. Traditional approaches to RBAC administration often are heavy weight in involving explicit actions by human administrators. These traditional approaches, where a human is in the loop in every administrative decision, are not scalable to the flexible and automated environment of knowledge management. Thus, one of the challenges for managing security in an environment is to make the administration as seamless as possible. The assignment of users and permissions to appropriate roles should occur transparently as part of the normal workflow of the organization. Recently, Sandhu and co-workers have introduced lightweight administration models for RBAC based on user attributes [1] and have also examined interaction of roles and workflow [13]. We need to build upon this work to develop administrative models for RBAC for knowledge management with the goal of being as lightweight and seamless as possible without compromising security. Traditional approaches to RBAC administration are also focused on the needs of a single organization or environment. This reflects the roots of this technology in enterprise access control. These approaches to security management do not scale to a cross-coalition global environment. Different local environments may use completely different roles or use the same role names but with very different meaning. The hierarchical relationships between the roles may be different

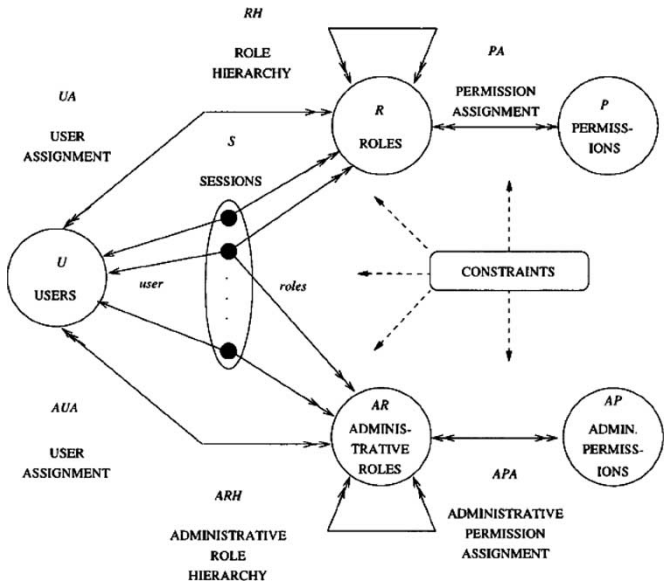

Fig. 3. RBAC model.

and possibly inconsistent. In some environments, a professor role may be senior to a student role, whereas in other environments within the same university, the student role may be senior to professor. For example, a student may be a dean taking a class from a professor, but this particular student is also the professor's boss. Another example would be a student who is the president of an association taking a class from a professor who is a member of the same association.

The basic structure of RBAC as developed in [18] is illustrated in Fig. 3. One of the key questions in applying RBAC to secure knowledge management is the nature of permissions in the knowledge-management context. The RBAC model is deliberately silent about the nature of permissions since this is highly dependent on the kind of system or application under consideration. Secure knowledge management requires control of access to a diverse set of information resources and services. We can identify the following broad categories:

1) information sources including structured and unstructured data, both within the organization and external to the organization;
2) search engines and tools for identifying relevant pieces of this information for a specific purpose;
3) knowledge extraction, fusion, and discovery programs and services;
4) controlled dissemination and sharing of newly produced knowledge.

Access to structured information in an organization, which resides in the organization's databases and other application repositories, is likely to already be under the purview of some form of RBAC using organizational roles. It is reasonable to assume that these organizational roles could be the foundation for role-based access to unstructured information.

However, access to unstructured information, which may reside on individual users' personal computers or small department-level servers, is fundamentally more problematic. Organizations will need to articulate and enforce access-control policies with respect to this information. While the initial thrust

of knowledge management has been on techniques to extract useful knowledge from this scattered but extremely valuable information, challenging access-control issues must be addressed before these techniques can be applied in production systems. We can assume a facility for distinguishing shared information from private information on a user's personal computer.

Existing personal-computer platforms have notoriously weak security so it will require fundamental enhancements in personal-computer technology to give us a reasonable level of assurance in this regard. Fortunately, there are several industry initiatives underway and, hopefully, some of these will come to fruition. The users will also need to determine how to share the public information. While reasonably fine-grained sharing techniques are available, it is unreasonable and undesirable to rely on the end users to specify these policies in detail. Moreover, current approaches are based on individual identities rather than roles. Scalability of information-sharing policies will require that they use organizational roles as a foundation. Information-sharing policies must be under control of the organization and cannot degenerate into anarchy where every user shares what they feel appropriate. There is a strong industry trend towards managing the corporate user's platform for a variety of reasons, so the requirement to impose role-based information-sharing policies would be consistent with this trend.

Access to specific search engines might involve access controls because of the cost or sensitivity issues. Cost comes about in terms of cost of licenses and such for accessing the search engine as well as the cost of the actual effort of performing the search. Sensitivity comes about in terms of sensitivity of the results obtained by the search. Similar comments apply to the knowledge-extraction algorithms. Finally, the resulting knowledge itself needs to be shared and protected, thus completing the cycle.

The challenge for RBAC is to go beyond the traditional picture of human administration of authorizations as depicted by the administrative roles in Fig. 3, and move to a more seamless and less user-intrusive administrative model. More generally, we may need to invent new forms of RBAC that allow users to do some degree of exploration in the information space of an organization without allowing carte blanche access to everything. This presents a significant research challenge to information-security researchers.

2) UCON for Knowledge Management: The concept of UCON was recently introduced in the literature by Park and Sandhu [17]. In recent years, there have been several attempts to extend access-control models beyond the basic access matrix model of Lampson, which has dominated this arena for over three decades. UCON unifies various extensions proposed in the literature in context of specific applications such as trust management and digital rights management. The UCON model provides a comprehensive framework for the next-generation access control. A UCON system consists of six components: subjects and their attributes, objects and their attributes, rights, authorizations, obligations, and conditions. The authorizations, obligations, and conditions are the components of the UCON decisions. An authorization rule permits or denies access of a subject to an object with a specific right based on the subject and/or object attributes, such as role name, security classification or clearance, credit amount, etc. An attribute is regarded as a variable with a value assigned to it in each system state. UCON is an attribute-based model, in which permission is authorized depending on the values of subject and object attributes. UCON extends the traditional access-control models in one aspect that the control decision depends not only on authorizations, but also on obligations and conditions. Obligations are activities that are performed by the subjects or by the system. For example, playing a licensed music file requires a user to click an advertisement and register in the author's web page. Such an action can be required before or during the playing process. Conditions are system and environmental restrictions that are not directly related to subject or object attributes, such as the system clock, the location, system load, system mode, etc. Another aspect that UCON extends traditional access-control models is the concepts of continuity and mutability.

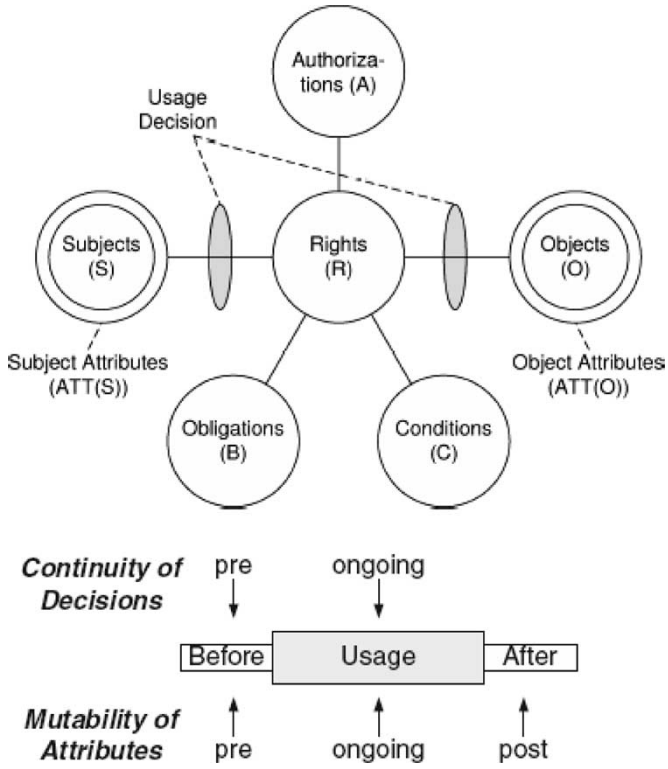

Fig. 4. UCON components.

A complete usage process consists of three phases along the time—before usage, ongoing usage, and after usage. The control-decision components are checked and enforced in the first two phases, named predecisions and ongoing decisions, respectively. The presence of ongoing decisions is called continuity, as indicated Fig. 4. Mutability means that the subject or object attribute value may be updated to a new value as a result of accessing. Along with the three phases, there are three kinds of updates: preupdates, ongoing updates, and postupdates. All these updates are performed and monitored by the security system as the access being attempted by the subject to the object. Updating of attributes by the side effect of subject activity is a significant extension to classic access control where the reference monitor mainly enforces existing permissions. Changing subject and object attributes has an impact on other ongoing

or future usage of permissions involving this subject or object. This aspect of mutability makes UCON very powerful. The new expressive power brought in by UCON is very germane to the automated and seamless security administration required in environments. The core UCON model in [17] is illustrated in Fig. 4.

The concept of a role is easily accommodated in UCON by means of attributes. A role is simply another attribute attached to subjects and objects. UCON extensions to the RBAC model will be very useful in specifying access-control policies for secure knowledge management. The authorization component of UCON is of course fundamental to access control. However, the concept of attribute mutability brings in an important additional component that allows us to control the extent of service provided to a user based on cost. Mutable attributes are automatically updated as a consequence of access. Thus, they can be used to limit the cost incurred by a knowledge search or knowledge extraction. Attribute mutability can also be used to control the scope of a knowledge search. Information about the past activities of a user can be aggregated in attributes associated with the user so these now become available for controlling future access to others.

The concept of obligations in UCON also is beneficial for knowledge management. Obligations are actions required to be performed before access is permitted. Obligations can be used to determine whether or not an expensive knowledge search is required or whether the knowledge is already available in the system. Obligations can also be used to escalate anomalous resource usage by a user so as to require additional conforming of the business need for the requested activity by that user or some other user. Ongoing obligations can be brought into play to regulate the amount of resources spent on a knowledge search or extraction operation when these resources may not be predictable in advance.

The use of conditions can allow resource-usage policies to be relaxed during times of low activity or tightened during peak periods of high usage. This can be correlated with cost attributes. Thus, the cost at times of high load may be higher than the cost at times of low load. UCON facilitates the automatic adaptation of such policies.

IV. Trust Management and Negotiation

Trust management and negotiation is a key aspect of secure knowledge management. Knowledge management is essentially about corporations sharing knowledge to get a competitive advantage. This means that one needs to trust the individuals with whom he or she is prepared to share the knowledge. Furthermore, corporations may have to negotiate contracts about knowledge sharing, and this means that a corporation has to trust another corporation before drawing up the contracts. This section will discuss trust management and negotiation concepts and issues applicable for knowledge management.

In today's web-based knowledge-intensive applications, trust is an important goal. A question that many users have when interacting with a web server, with an application, or with an information source is "Can I trust this entity?" Ideally, one would like to have available tools able to automatically evaluate the level of trust one may place on an unknown party encountered on the web. Though we are still far from achieving such a goal, the problem of trust is currently being actively investigated and several notions and tools addressing aspects related to trust have been developed. Trust management and negotiation is becoming an important aspect of secure knowledge management. When knowledge is shared across and within organizations, the parties involved have to establish trust rules for collaboration. Therefore, trust management plays an important role in knowledge management.

In this section, we first discuss possible definitions of trust to give an idea of the different meanings associated with it, and we point out the notion that we refer to in the remainder of the discussion. We then discuss two different issues related to trust: how trust is established and how it is managed. Finally, we discuss a class of systems, known as trust-negotiation (TN) systems, that establishes a form of trust through a controlled exchange of credentials among some interacting parties. It is important to notice that trust, in its many forms, relies on information and knowledge about the interacting entities. As such, TN is an area where sophisticated knowledge-management tools could be profitably used. Furthermore, trust management is a key security technique for knowledge management. This is because organizations have to establish trust before sharing data, information, and knowledge.

1) Some Definitions of Trust: The notion of trust is used in a large number of different contexts and with diverse meanings, depending on how it is used. It is a complex notion about which no consensus exists in the computer and information-science literature, although its importance has been widely recognized. Different definitions are possible depending on the adopted perspective. For example, Kini and Choobineh [14] define trust from the perspectives of personality theorists, sociologists, economists, and social psychologists. They highlight the implications of these definitions and combine their results to create their definition of trust in a system. They define trust as: "a belief that is influenced by the individual's opinion about certain critical system features." Their analysis covers various aspects of human trust in computer-dependent systems, but they do not address the issue of trust between parties (humans or processes) involved in web-based transactions.

A different definition is based on the notion of competence and predictability. The Trust-EC project (http://dsa-isis.jrc.it/TrustEC/) of the European Commission Joint Research Centre (ECJRC) defines trust as: "the property of a business relationship, such that reliance can be placed on the business partners and the business transactions developed with them." Such definition emphasizes the identification and reliability of business partners, the confidentiality of sensitive information, the integrity of valuable information, the prevention of unauthorized copying and use of information, the guaranteed quality of digital goods, the availability of critical information, the management of risks to critical information, and the dependability of computer services and systems. Another relevant definition is by Grandison and Sloman [10], and they define trust as: "the firm belief in the competence of an entity to act dependably, securely, and reliably within a specified context." They argue that

trust is a composition of many different attributes—reliability, dependability, honesty, truthfulness, security, competence, and timeliness—which may have to be considered depending on the environment in which trust is being specified.

A main difficulty of all these definitions as well as others is that they provide a notion of trust for which establishing metrics and developing evaluation methodologies are quite difficult. We thus adopt a more restricted notion of trust, which is the one underlying TN systems. Such notion was initially proposed by Blaze and Feigenbaum, according to whom, "Trust management problems include formulating security policies and security credentials, determining whether particular sets of credentials satisfy the relevant policies, and deferring trust to third parties" [5]. Such a definition of trust basically refers to security policies regulating accesses to resources and credentials that are required to satisfy such policies. TN thus refers to the process of credential exchanges that allows a party requiring a service or a resource from another party to provide the necessary credentials in order to obtain the service or the resource. Notice that, because credentials may contain sensitive information, the party requiring the service or the resource may ask to verify the other party's credentials before releasing its own credentials. This definition of trust is very natural for secure knowledge management as organizations may have to exchange credentials before sharing knowledge.

2) Trust Services: The notion of trust services is not new; there are various financial, insurance, and legal services available that make business activities simpler and less risky. In the web-based transaction context, trust services are emerging as a business enabler, with the goal of delivering trust and confidence at various stages of the interaction among the parties involved in a transaction, including: establishing and maintaining trust, negotiations, contract formation, fulfillment, collaboration, through to dispute resolution. Trust services attempt to solve problems such as: establishing the authenticity of electronic communications; ensuring that electronic signatures are fair and legally binding, and creating an electronic audit trail that can be used for dispute resolution. The area of trust services is today a fast-moving area and it is difficult to anticipate the range of trust services that will be available in the next few years. We can, however, reasonably expect that they include mechanisms to support trust establishment, negotiation, agreement, and fulfillment, such as identity services, authorization services, and reputation services. Knowledge-management strategies make use of the trust services.

3) TN Systems: TN is an emerging approach exploiting the concept of properties of the entities as a means for establishing trust, particularly in open environments such as the web, where interacting entities are usually unknown to each other. TN is a peer-to-peer interaction, and consists of the iterative disclosure of digital credentials, representing statements certified by given entities, for verifying properties of their holders in order to establish mutual trust. In such an approach, access to resources (data and/or services) is possible only after a successful TN is completed. A TN system typically exploits digital identity information for the purpose of providing a fine-grained access control to protected resources. However, unlike conventional access-control models, TN assumes that the interacting parties

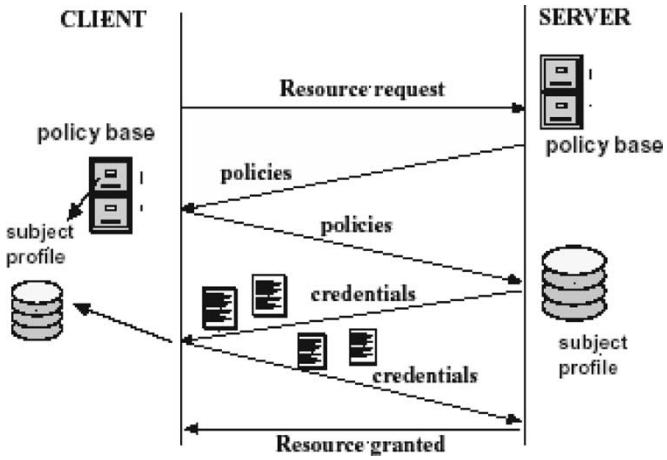

Fig. 5. Organization of a TN process.

are peer and that each peer needs to be adequately protected. For instance, with respect to the peer owning the resource to be accessed, assets that need to protected are, in addition to the resource, the access-control policies, as they may contain sensitive information, and the credentials of the resource owner. With respect to the peer requiring access to the resource, the assets to be protected are the credentials as they often contain private information about the individual on behalf of whom the peer is negotiating.

4) TN Building Blocks: We now briefly describe how negotiations are generally intended, and identify the main phases and functional components of negotiations as given in [2]. Such an approach is ideal for knowledge management. A TN involves two entities, namely a client, which is the entity asking for a certain resource, and a server, which is the entity owning (or more generally, managing access to) the requested resource. The model is peer to peer: Both entities may possess sensitive resources to be protected, and thus must be equipped with a compliant negotiation system. The notion of resource comprises both sensitive information and services, whereas the notion of entity includes users, processes, roles, and servers. The term resource is intentionally left generic to emphasize the fact that the negotiations we refer to are general purpose, that is, a resource is any sensitive object (e.g., financial information, health records, and credit card numbers) whose disclosure is protected by a set of policies.

Fig. 5 illustrates a typical negotiation process. During the negotiation, trust is incrementally built by iteratively disclosing digital credentials in order to verify properties of the negotiating parties. Credentials are typically collected by each party in appropriate repositories, called subject profiles. Another key component of any TN is a set of access-control policies, referred to as disclosure policies, governing access to protected resources through the specification of the credential combinations that must be submitted to obtain access to the resources.

To carry out a TN, parties usually adopt a strategy, which is implemented by an algorithm determining which credentials to disclose, when to disclose them, and whether to succeed or fail the negotiation. Several TN strategies can be devised, each with different properties with respect to speed of negotiations and caution in releasing credentials and policies. The efficiency of

a strategy depends on two factors: the communication cost and the computational cost. The communication cost includes the sizes of the messages exchanged and their number. Communication and computational costs of a negotiation strictly depend on the adopted strategy and vary from exponential, when a brute-force strategy is adopted, to more efficient strategies.

5) TN Requirements: We now discuss the relevant dimensions with respect to which policy languages can be analyzed. These dimensions can be classified in two main groups, i.e., those related to the adopted language and those related to the system and its components. It is important to note that these requirements are a partial list and other requirements are likely to be identified as research and deployment of negotiation systems progress, given also the increasing number of researchers actively contributing to the trust-management area.

a) Language requirements: TN policy languages are a set of syntactic constructs (e.g., credentials, policies) and their associated semantics, encoding security information to be exchanged during negotiations. Effective TN languages should be able to simplify credential specification and also to express a wide range of protection requirements through specification of flexible disclosure policies. The main relevant dimensions for these languages are related with expressiveness and semantics.

b) System requirement: The development of comprehensive TN systems is quite challenging. On the one hand, such systems should be flexible, scalable, and portable. On the other, they should support advanced functions, such as support for credential chains, authentication of multiple identities, and complex compliance-checking modes whose efficient implementation is often difficult. In particular, the compliance checker should be able to interpret a remote policy and check whether there exists a set of local credentials that satisfy the received policy.

6) Selected TN Systems: Because of the relevance of TN for web-based applications and knowledge management, several systems and research prototypes have been developed. The most well-known systems include KeyNote by Blaze *et al.* [5], TrustBuilder by Yu and Winslett [26], and Trust-X by Bertino *et al.* [3], which we briefly discuss in what follows. These systems are relevant for TN in knowledge-management systems.

KeyNote has been developed to work for large- and small-scale Internet-based applications. It provides a single unified language for both local policies and credentials. KeyNote credentials, called assertions, contain predicates that describe the trusted actions permitted by the holders of a specific public key. As a result, KeyNote policies do not handle credentials as a means of establishing trust, mainly because the language was intended for delegation authority. Therefore, it has several shortcomings with respect to the requirements we have outlined.

TrustBuilder is one of the most significant systems in the negotiation research area. It provides several negotiation strategies, as well as a strategy- and language-independent negotiation protocol ensuring the interoperability of the defined strategies. In TrustBuilder, each negotiation participant has an associated security agent that manages the negotiation. During a negotiation, the security agent uses a local negotiation strategy to determine which local resources to disclose next and to accept new disclosures from other parties. TrustBuilder includes a credential-verification module, a policy-compliance checker, and a negotiation-strategy module, which is the system's core. The system relies on a credential-verification module, which performs a validity check of the received credentials. Some recent investigation carried out in the framework of TrustBuilder includes support for sensitive policies (obtained by introducing hierarchies in policy definitions), and privacy protection mechanisms (obtained by introducing dynamic policies, that is, policies dynamically modified during a negotiation).

Trust-X supports all aspects of negotiation, specifically developed for peer-to-peer environments. Trust-X supports an Extensible Markup Language (XML)-based language, known as XML-based Trust Negotiation Language (X-TNL), for specifying Trust-X certificates and policies. Trust-X has a typing credential system and addresses the issue of vocabulary agreement using XML namespaces. The use of namespaces combined with the certificate-type system helps the TN software in correctly interpreting different credentials' schema, even when issued by different entities not sharing a common ontology. A novel aspect of X-TNL is its support for special certificates, called trust tickets. Trust tickets are issued on successfully completing a negotiation and can speed up subsequent negotiations for the same resource. X-TNL provides a flexible language for specifying policies and a mechanism for policy protection, based on the notion of policy preconditions. A Trust-X negotiation consists of a set of phases that are sequentially executed. In particular, Trust-X enforces a strict separation between policy exchange and resource disclosure. This distinction results in an effective protection of all the resources involved in negotiations. Trust-X is a flexible system, providing various TN strategies that allow better tradeoffs between efficiency and protection requirements. In particular, Trust-X supports three different negotiation modes. The first, based on trust tickets, can be adopted when the parties have already successfully completed a negotiation for the same resource. The second mode, based on using specific abstract data structures called negotiation trees, performs a runtime evaluation of the negotiation's feasibility by determining a sequence of certificate disclosures that can successfully end the negotiation. The last mode exploits a notion of similarity between negotiations and is based on the observation that a service provider usually handles many similar negotiations. Some recent investigation carried out in the framework of Trust-X includes support for privacy, anonymity, the development of a recovery protocols, and the integration with federated identity management.

V. Privacy Management

Secure knowledge-management technologies include technologies for secure data management and information management including databases, information systems, semantic web, and data mining. For example, data mining is an important tool in making the web more intelligent. Because of its ability to extract data on the web, it can aid in the creation of the semantic web and, subsequently, in the knowledge-management strategies and processes. The semantic web can be utilized

by the knowledge manager to execute the various strategies and processes as well as to collect metrics. However, both the semantic web and data mining have inferencing capabilities, and therefore, one could combine pieces of information together and infer information that is highly private or highly sensitive [20]. This problem is the privacy problem [21].

Now, the semantic-web community has come up with platform for privacy preferences (P3P) [16] specifications. That is, when a user enters a web site, the site will specify its privacy policies and the user can then submit information if he/she desires. However, the web site may give out say medical records and names separately to a third party such as an advertising agency. This third party can make unauthorized inferences from the information legitimately obtained from the web site and deduce private information that associates the medical records with the names. That is, while data mining and the semantic web are very useful knowledge-management technologies, for secure knowledge management, we need to ensure that the information extracted as a result of data mining does not compromise security and privacy. Furthermore, the semantic-web engine has to go beyond P3P enforcement to ensure privacy. In this section, we will discuss an approach to ensuring privacy for the semantic web so that privacy is managed as part of the knowledge-management process.

Since data may be mined and patterns and trends extracted, security and privacy constraints can be used to determine which patterns are private and sensitive and to what extent. For example, suppose one could extract the names of patients and their corresponding healthcare records. If a privacy constraint states that names and healthcare records taken together are private, then this information is not released to the general public. If the information is semiprivate, then it may be released to those who have a need to know, such as say a healthcare administrator. Essentially, the inference-controller approach discussed in [19] is one solution to achieve some level of privacy. That is, the inference controller examines the privacy constraints, the query, and the information that has been released, and determines whether to release any further information. This approach can be regarded to be a type of privacy-sensitive data mining [22]. In our research, we have found many challenges to the inference-controller approach as discussed in [19]. For example, how long can the system maintain the history information? How can we ensure that the constraints are consistent and complete? These challenges will have to be addressed when handling security and privacy constraints for the semantic web. Fig. 6 illustrates security/privacy controllers for the semantic web. As illustrated, there are data-mining tools on the web that mine the web databases. The privacy controller should ensure privacy-preserving data mining. Ontologies may be used by the privacy controllers. For example, there may be ontology specifications for privacy constraints and these specifications may be used by the inference controller to reason about the applications. Furthermore, XML and resource description framework (RDF) may be extended to specify security and privacy policies [4], [6].

The secure knowledge manager will utilize the secure semantic web to execute its knowledge-management strategies and processes. For example, the semantic web will be utilized

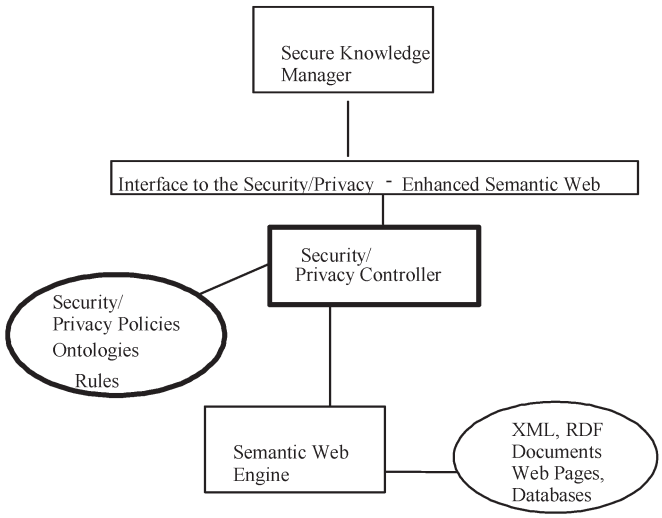

Fig. 6. Privacy controller for the semantic web.

for various operations such as order management, procurement, and contracting. The security and privacy controllers will ensure that the security and privacy rules are not violated when executing the knowledge-management processes. We have discussed a high-level design. The next step is to design the details of the security and privacy controllers and implement a secure semantic web for secure knowledge management. Some details of the design and algorithms are given in [24] and [23].

VI. Summary and Directions

This paper has discussed some key points in secure knowledge management. We have stressed that security has to be incorporated into the knowledge-management lifecycle. We discussed issues on integrating security strategy with knowledge management and business strategies of an organization. We also discussed an architecture for secure knowledge management. Then, we focused on two prominent security techniques. With respect to access control, we discussed both RBAC and UCON and showed with examples how these approaches may be applied for knowledge management. Then, we discussed trust management and negotiation for knowledge management. Because knowledge management will involve multiple organizations or multiple departments within an organization, it is very important that the different parties establish TN rules for collaboration. Finally, we discussed privacy for secure knowledge management including privacy problems that arise due to data mining and inferencing inherent to the semantic web.

There are many areas that need further work. First, we need to develop a methodology for secure knowledge management. While we have discussed some aspects of secure knowledge-management strategies, processes, and metrics, we need a comprehensive lifecycle for secure knowledge management. We also need to investigate further RBAC and UCON, as well as trust management and negotiation. The definitions and rules discussed in this paper have to be formalized for RBAC, UCON, and trust. The security critical components have to be

identified for knowledge management. Finally, privacy issues need to be investigated further.

In addition to enhancing and formalizing the policies discussed here, we also need to explore the incorporation of some of the real-world policy specifications into the knowledge-management strategies. For example, we need to examine the P3P specified by the World Wide Web Consortium and determine how we can enforce such a policy within the framework of secure knowledge management. We also need to investigate integrity aspects of knowledge management. For example, how do we ensure the integrity of the data and the activities? How can we ensure data, information, and knowledge quality? The best way to test out the policies is to carry out pilot projects for different types of organizations including those from academia, industry, and government. Based on the results obtained, we can then continue to refine the policies for knowledge management.

In summary, secure knowledge management will continue to be critical as organizations work together, share data, as well as collaborate on projects. Protecting the information and activities while sharing and collaborating will be a major consideration. This paper has provided some directions for incorporating confidentiality, trust, and privacy for knowledge management.

ACKNOWLEDGMENT

The authors would like to thank N. Tsybulnik and L. Liu for comments on this paper and their research on secure semantic web and privacy-preserving data mining.

REFERENCES

[1] M. Al-Kahtani and R. Sandhu, "A model for attribute-based user-role assignment," in *Proc. 17th Annu. Computer Security Applications Conf.*, Las Vegas, NV, Dec. 2002, pp. 353–362.

[2] E. Bertino, E. Ferrari, and A. C. Squicciarini, "Trust negotiations: Concepts, systems and languages," *Comput. Sci. Eng.*, vol. 6, no. 4, pp. 27–34, Jul./Aug. 2004.

[3] ——, "A peer-to-peer framework for trust establishment," *IEEE Trans. Knowl. Data Eng.*, vol. 16, no. 7, pp. 827–842, Jul. 2004.

[4] E. Bertino, B. Carminati, E. Ferrari, and B. Thuraisingham, "Secure third party publication of XML documents," *IEEE Trans. Knowl. Data Eng.*, vol. 16, no. 10, pp. 1263–1278, Oct. 2004.

[5] M. Blaze, J. Feigenbaum, and J. Lacy, "Decentralized trust management," in *Proc. IEEE Symp. Security Privacy*, 1996, pp. 164–173.

[6] B. Carminati, E. Ferrari, and B. Thuraisingham, "Securing RDF documents," in *Proc. DEXA Workshop Web Semantics*, Zaragoza, Spain, Aug. 2004, pp. 472–480.

[7] G. Denker, L. Kagal, T. Finin, M. Paolucci, and K. Sycara, "Security for DAML web services: Annotation and matchmaking," in *Proc. Int. Semantic Web Conf.*, 2003, pp. 335–350.

[8] D. F. Ferraiolo, R. Sandhu, S. Gavrila, D. R. Kuhn, and R. Chandramouli, "Proposed NIST standard for role-based access control," *ACM Trans. Inf. Syst. Secur.*, vol. 4, no. 3, pp. 224–274, Aug. 2001.

[9] T. Finin and A. Joshi, "Agents, trust, and information access on the semantic web," *ACM SIGMOD Rec.*, vol. 31, no. 4, pp. 30–35, Dec. 2002.

[10] T. Grandison and M. Sloman, "A survey of trust in Internet applications," *Commun. Surveys Tuts.*, vol. 3, no. 4, pp. 2–16, 4th Quarter 2000.

[11] L. Kagal, T. Finin, and A. Joshi, "A policy based approach to security for the semantic web," in *Proc. Int. Semantic Web Conf.*, 2003, pp. 402–418.

[12] L. Kagal, M. Paolucci, N. DSrinivasan, G. Denker, T. Finin, and K. Sycara, "Authorization and privacy for semantic web services," *IEEE Intell. Syst.*, vol. 19, no. 4, pp. 50–56, Jul./Aug. 2004.

[13] S. Kandala and R. Sandhu, "Secure role-based workflow models," in *Database Security XV: Status and Prospects*, D. Spooner, Ed. Norwell, MA: Kluwer, 2002.

[14] A. Kini and J. Choobineh, "Trust in electronic commerce: Definition and theoretical consideration," in *Proc. IEEE 31st Int. Conf. System Sciences*, 1998, pp. 51–61.

[15] D. Morey, M. Maybury, and B. Thuraisingham, Eds., *Knowledge Management*, Cambridge, MA: MIT Press, 2001.

[16] *Platform for Privacy Preferences*. [Online]. Available: www.w3c.org

[17] J. Park and R. Sandhu, "The UCONABC usage control model," *ACM Trans. Inf. Syst. Secur.*, vol. 7, no. 1, pp. 128–174, Feb. 2004.

[18] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman, "Role-based access control models," *Computer*, vol. 29, no. 2, pp. 38–47, Feb. 1996.

[19] B. Thuraisingham and W. Ford, "Security constraint processing in a multilevel secure distributed database management system," *IEEE Trans. Knowl. Data Eng.*, vol. 7, no. 2, pp. 274–293, Apr. 1995.

[20] B. Thuraisingham, "Security standards for the semantic web," *Comput. Stand. Interfaces*, vol. 27, no. 3, pp. 257–268, Mar. 2005.

[21] ——, *Database and Applications Security: Integrating Data Management and Information Security*. Boca Raton, FL: CRC Press, May 2005.

[22] ——, "Privacy preserving data mining: Developments and directions," *J. Database Manage.*, vol. 16, no. 1, pp. 75–87, Mar. 2005.

[23] ——, *Building Trustworthy Semantic Webs*. Boca Raton, FL: CRC Press, Jun. 2006.

[24] N. Tsybulnik, B. Thuraisingham, and L. Khan, "Design and implementation of a security controller for the semantic web," Univ. Texas, Dallas, Tech. Rep. UTDCS-06-06, 2005.

[25] S. Upadhyaya, H. R. Rao, and G. Padmanabhan, "Secure knowledge management," in *Encyclopedia of Knowledge Management*, D. Schwartz, Ed. Hershey, PA: Idea Group Inc., 2005.

[26] T. Yu and M. Winslett, "A unified scheme for resource protection in automated trust negotiation," in *Proc. IEEE Symp. Security and Privacy*, 2003, pp. 110–122.

Elisa Bertino (SM'83–F'02) received the B.S. and Ph.D. degrees in computer science from the University of Pisa, Pisa, Italy, in 1977 and 1980, respectively.

She is a Professor of computer science and electrical and computer engineering at Purdue University, West Lafayette, IN, and serves as Research Director of CERIAS. Previously, she was a Faculty Member at the Department of Computer Science and Communication of the University of Milan where she directed the DB&SEC laboratory. She has been a Visiting Researcher at the IBM Research Laboratory (now Almaden) in San Jose, at the Microelectronics and Computer Technology Corporation, at Rutgers University, and at Telcordia Technologies. She has been a consultant to several Italian companies on data management systems and applications and has given several courses to industries. She has been involved in several projects sponsored by the EU. Her main research interests include security, privacy, database systems, object-oriented technology, multimedia systems. In those areas, she has published more than 250 papers in all major refereed journals, and in proceedings of international conferences and symposia. She a Coeditor in Chief of the *Very Large Database Systems (VLDB) Journal*. She is a coauthor of the books *Object-Oriented Database Systems—Concepts and Architectures* (Addison-Wesley, 1993), *Indexing Techniques for Advanced Database Systems* (Kluwer, 1997), and *Intelligent Database Systems* (Addison-Wesley, 2001). She also serves on the editorial boards of several scientific journals, incuding *IEEE Internet Computing*, *ACM Transactions on Information and System Security*, *Acta Informatica*, *the Parallel and Distributed Database Journal*, *the Journal of Computer Security*, *Data and Knowledge Engineering*, *the International Journal of Cooperative Information Systems*, and *Science of Computer Programming*.

Prof. Bertino is a member of the advisory board of the IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING. She has served as Program Co-Chair of the 1998 IEEE International Conference on Data Engineering (ICDE). She has served as Program Committee members of several international conferences, such as ACM SIGMOD, VLDB, and ACM OOPSLA, as Program Chair of the 2000 European Conference on Object-Oriented Programming (ECOOP 2000), and as Program Chair of the 7th ACM Symposium of Access Control Models and Technologies (SACMAT 2002). She recently served as Program Chair of the 2004 EDBT Conference. She is a Fellow of the ACM and has been named a Golden Core Member for her service to the IEEE Computer Society. She received the 2002 IEEE Computer Society Technical Achievement Award for "outstanding contributions to database systems and database security and advanced data management systems."

Latifur R. Khan received the B.Sc. degree in computer science and engineering from the Bangladesh University of Engineering and Technology, Dhaka, in 1993, and the M.S. and Ph.D. degrees in computer science from the University of Southern California, Los Angeles, in 1996 and 2000, respectively.

He has been an Assistant Professor of Computer Science department at the University of Texas at Dallas since September 2000. He is currently supported by Grants from the National Science Foundation, Nokia Research Center, Texas Instruments, and Alcatel, USA. He has authored more than 70 articles, book chapters, and conference papers focusing in the areas of multimedia information management, data mining, and intrusion detection. He has also served as a Referee for database/data mining journals and conferences (e.g., IEEE TKDE, KAIS, ADL, VLDB), and he is currently serving as a Program Committee Member for the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD2005), International Conference on Database and Expert Systems Applications DEXA 2005, and International Conference on Cooperative Information Systems (CoopIS 2005), and Program Co-Chair of the ACM SIGKDD International Workshop on Multimedia Data Mining, 2005. He has also given a tutorial in the 14th ACM International World Wide Web Conference, WWW2005, May 2005, Chiba, Japan, and he has been an Associate Editor of *Computer Standards and Interfaces Journal* by Elsevier Publishing since June 2005.

Dr. Khan has been awarded the Sun Equipment Grant.

Ravi Sandhu (M'89–SM'90–F'02) received the B.Tech. and M.Tech. degrees in electrical engineering from the Indian Institutes of Technology at Bombay and Delhi, respectively, and the M.S. and Ph.D. degrees in Computer Science from Rutgers University, NJ, in 1975, 1980, and 1983, respectively.

He is a Professor of Information and Software Engineering and Director of the Laboratory for Information Security Technology at the George Mason University, Fairfax, VA. He is a leading authority on access control, authorization, and authentication models and protocols. His seminal paper on role-based access control (RBAC) introduced the RBAC96 model, which evolved into the 2004 National Institute of Standards and Technology/American National Standards Institute (NIST/ANSI) standard RBAC model (and is on track to become an ISO standard). More recently, he introduced the usage control (UCON) model as a foundation for next-generation access control by integrating obligations and conditions with the usual notion of authorization in access control and providing for continuity of enforcement and mutability of attributes. He has also served as the principal designer and security architect of TriCipher's Armored Credential System (TACS), which earned the coveted FIPS 140 level-2 rating from NIST. He has provided high-level security consulting services to several private and government organizations. His research has been sponsored by numerous public and private organizations currently including Lockheed Martin, Northrop Grumman, Intel, Verizon, Network Associates, Defense Advanced Research Projects Agency, Department of Defense, Department of Energy, National Security Agency, National Reconnaissance Agency, National Science Foundation, Naval Research Laboratory, Internal Revenue Service, and the Advanced Research and Development Activity agency. Previously, he has published influential and widely cited papers on various security topics including safety and expressive power of access-control models, lattice-based access controls, and multilevel secure relational and object-oriented databases. He has published over 160 technical papers on computer security in refereed journals, conference proceedings and books. He founded the *ACM Transactions on Information and Systems Security (TISSEC)* in 1997 and served as Editor-In-Chief until 2004. He served as Chairman of ACM's Special Interest Group on Security Audit and Control (SIGSAC) from 1995 to 2003, and founded and led the ACM Conference on Computer and Communications Security (CCS) and the ACM Symposium on Access Control Models and Technologies (SACMAT) to high reputation and prestige.

Dr. Sandhu is a Fellow of the ACM. Most recently, he founded the IEEE Workshop on Pervasive Computing Security (PERSEC) in 2004.

Bhavani Thuraisingham (SM'97–F'03) received the B.S. degree in mathematics and physics from the University of Ceylon, Sri Lanka, in 1975, the M.S. degree in mathematics from the University of Bristol, Bristol, U.K., in 1977, and the Ph.D. degree in computer science from the University of Wales, Wales, U.K., in 1979.

She has served as an expert consultant in information security and data management to the Department of Defense, the Department of Treasury, and the Intelligence Community for over ten years and is an Instructor for Armed Forces Communication and Electronics Association (AFCEA) since 1998. Prior to joining the University of Texas at Dallas, she was an Intergovernmental Personnel Act (IPA) at the National Science Foundation (NSF) from MITRE Corporation. At NSF, she established the Data and Applications Security Program and cofounded the Cyber Trust theme and was involved in interagency activities in data mining for counterterrorism. She has been at MITRE from January 1989 until June 2005, and has worked in MITRE's Information Security Center and was later a Department Head in data and information management as well as Chief Scientist in data management. She has recently joined the University of Texas at Dallas as a Professor of Computer Science and Director of the Cyber Security Research Center in the Erik Jonsson School of Engineering. Her industry experience includes six years of product design and development of CDCNET at Control Data Corporation and research, development and technology transfer at Honeywell Inc. Her academia experience includes being a Visiting Faculty at the New Mexico Institute of Technology, Adjunct Professor of Computer Science first at the University of Minnesota and later at Boston University. Her research in information security and information management has resulted in over 70 journal articles, over 200 refereed conference papers, and three U.S. patents. She is the author of seven books in data management, data mining, and data security, including one on data mining for counterterrorism and another on database and applications security. She has given over 25 keynote presentations at various research conferences and has also given invited talks at the White House Office of Science and Technology Policy and at the United Nations on Data Mining for counter-terrorism. She serves (or has served) on editorial boards of top research journals. She is also establishing the consulting company "BMT Security Consulting" specializing in data and applications security consulting and training and is the Founding President of the company.

Ms. Thuraisingham is a Fellow of the American Association for the Advancement of Science (AAAS). She received IEEE Computer Society's prestigious 1997 Technical Achievement Award for "outstanding and innovative contributions to secure data management." She was elected as Fellow of the British Computer Society in February 2005.