RC24759 (W0903-025) March 6, 2009
Computer Science

IBM Research Report

A New Schema for Security in Dynamic Uncertain Environments

Dakshi Agrawal
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

LIMITED DISTRIBUTION NOTICE: This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties). Copies may be requested from IBM T. J. Watson Research Center , P. O. Box 218, Yorktown Heights, NY 10598 USA (email: reports@us.ibm.com). Some reports are available on the internet at http://domino.watson.ibm.com/library/CyberDig.nsf/home .

A New Schema for Security in Dynamic Uncertain Environments

Dakshi Agrawal
IBM T. J. Watson Research Center
agrawal@us.ibm.com

Abstract

It is our hypothesis that for a complex system of systems operating in a dynamic, uncertain environment the traditional approach of forward, static security is insufficient. What is required are macroscopic schemata for security that incorporate mechanisms which monitor the overall environment and feed their observations back into the security mechanisms so that they can adjust their 'posture' accordingly. Such schemata must also account for system-wide aggregated security risks in addition to risk presented by the individual users and information objects. We propose one such schema in this work.

To illustrate the utility of macroscopic schemata, we use the examples of two recent studies of access control systems and map their results to the proposed schema and distill macroscopic insights that are otherwise lost in details.

We hope that such security schemata will lead to a systematic analysis of security of complex systems akin to what is already available for complex social, biological, and mechanical systems. We hope that macroscopic models based on such schemata will be able to provide, through analysis, large-scale simulations, or by other means, a quantified assessment of the resilience of the security of a system of systems, and in the long run, provide systematic controls that can be used to adjust the security posture of a complex system.

1. Introduction

It is the uncertainty and dynamicity in the operating environment that ask the most penetrating questions from the current security solutions. In mechanical, chemical, and a host of other engineering systems, autonomic control mechanisms are used to compensate for uncertainty and dynamicity in the operating environment, however, few, if any, such mechanisms are explicitly designed in modern IT security systems. In the parlance of Control Theory, most security mechanisms are *open loop* since they do not *feedback* the output of their processes back into the security mechanisms to further regulate processes and drive their outputs to a desired state.

In the early days, computers were largely isolated from each other, had limited software functionality, and their users were technically sophisticated, resulting in an environment that was well-controlled. The properties of a security mechanism could be proved under a 'clean room' security-model that was not too far from the reality. However, computers have since then morphed into computing devices of all shapes and sizes; these devices are now connected together resulting in networks of all kinds and reaches; the functionality has grown exponentially and the user base has expanded to include technical equivalent of laity. While the current state of the art in computer security has addressed many challenges rising from these changes, it has failed to systematically address the most basic change; namely, there is a lot more uncertainty and dynamism in the operating environment and the context of computing systems today than it was a few decades ago.

It is our belief that to provide security for a complex system of systems, new schemata for security are needed that explicitly incorporate mechanisms which monitor the overall environment and feed their observations back into the security mechanisms to compensate for uncertainty and dynamicity in the environment. It also follows that individual security mechanisms need to be designed so that they provide 'control' parameters that can be set to adjust security properties of the system to keep them within acceptable limits even as the environment and the context of the system undergoes changes. We need to enhance our understanding of systematic properties of security systems, and ultimately, provide a global *dashboard* for security. Such desired advances are complementary to the traditional approach of investigating properties of carefully designed security mechanisms under sterile security models: they will fill critical gaps in current security solutions which have been exposed by the diversity of goals and needs of the users, and above all, the increasingly networked nature of interdependent social, information, and communication infrastructures.

The paper is organized as follows. In Section 2, we present an example security schema for access control that

addresses many of the questions raised above. With the help of two examples, we then explore how the proposed schema can be used to provide interesting insights into the design of access-control systems. We briefly comment on feasibility of implementation of the system before concluding the paper in Section 3.

2. A New Schema for Information Sharing

Figure 1 shows a simplified macroscopic schema of information sharing from the security perspective. The goal of this simplified schema is to illustrate a few key aspects of the complex processes that go towards making information sharing decisions and to highlight the interconnectedness of these processes.

The schema in Figure 1 consider a scenario in which a user U wants access to an information object O. We leave the details of the access control system A unspecified at the moment, and assume that it is a risk-benefit aware access control system [7] which allows accesses based on a variety of inputs including the context of an access request. We assume that the access control system is flexible, and it provides systematic control parameters that can be easily set to tune the computation of risks and benefits associated with a request to reflect dynamic, uncertain operating environment. If the user U is granted access to the object O, then there is a chance that the object will be leaked by the user, or be misused otherwise, causing harm to the organization. This leakage or misused is represented in Figure 1 by a "transfer function" L, and the consequent organizational damage is denoted by Q. In a simple case, input to L is a triplet of the form $\langle O, U, A(O, U)\rangle$ and the output is a binary variable that represents whether the object O has been leaked. In this schema, we assume that a leaked object is available to everyone. We also assume that the organizational damage Q depends on the aggregated set of all leaked objects.

We assume that the leakage of information objects is monitored by a monitor M which may be imperfect. The output of the monitoring is fed back to the access control subsystem so that the past leakage history can be taken into account while making a decision on the future access requests by the user U, or for the object O. We assume that the user U derives a personal benefit B by accessing the object O while incurring a cost for access. This cost depends on attributes of the user U and the object O, and in addition, on past monitoring results. The net benefit of an access to the user is a function N of the personal benefit B and the personal cost C.

We assume that information access also results in benefits for the organization. In Figure 1, organizational benefits are denoted by P and they are a function of the aggregate of all information accesses. The goal of the organization is to maximize the organizational net benefit G which is a function of organizational benefits and damages subjected to any constraint that the organization may face. An individual's net benefit may be influenced by the overall organizational benefit G.

2.1. Using the Schema

Let us first list major assumptions used in sketching the security schema given above. The first major assumption is that it is feasible to monitor leakage of information. Second, we assume that users can explicitly estimate 'benefits' of an access, and it is possible to estimate organizational benefits and damages. Third, we assume that there is a mechanism in place so that users can be charged a cost for accessing information. It is implicit in our schema that these functions be performed without significant human involvement. Finally, to exploit the monitoring information and other contextual and environmental inputs, we assume that there is a flexible access control mechanism that can be tuned dynamically and with ease.

Satisfying any of the assumptions listed above would require significant innovation in the state of the art in security mechanisms. Therefore, it is natural to ask the question in which direction should the efforts be directed; in other words, how sensitive is the security system represented by the given schema to different parameters that characterize individual subsystems. How can the overall design be analyzed, viz., can the benefits of a configuration be quantified in terms of increase in information flow, reduction in risk, reduction in security exception grants, etc. so as to optimize the system design.

To give an example, the timescale and granularity of various processes have been left unspecified in the schema given above. A monitor may estimate leakage at different time scales with different granularity: it may monitor leakage per individual user per access, for the whole organization over an extended duration, or at a granularity that is in between, for example, by roles per week. Experience from Industrial Engineering informs us that organizational benefits and damages are likely to be estimated coarsely over a relatively longer time period. As a result, different processes in the schema use time epochs of different lengths that may not be synchronized with each other. How does that affect the efficacy of the overall design?

We believe that new macroscopic security schemata are needed to answer such question. A diverse set of techniques, from fields as diverse as control theory, variational calculus, agent-based simulations, etc., that have been successfully applied in Economics, Traffic Engineering, Ecology, Social Sciences, etc., then can be applied to security design. In the following, we give two examples that illustrate the utility of analyzing security in terms of macro-

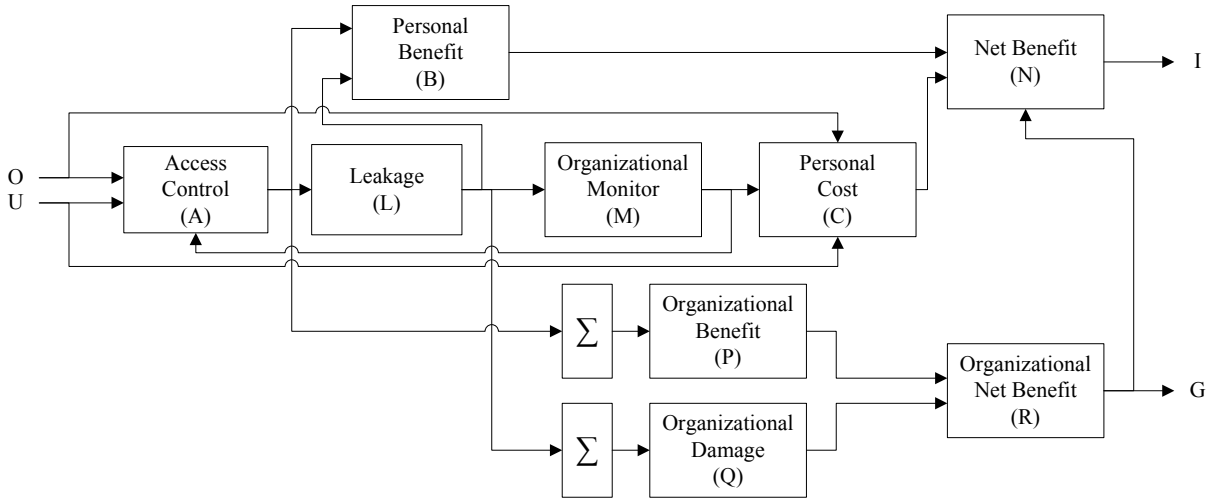

Figure 1. An Schema of an Access Control System

scopic schemata. These examples have been decocted from two recent results—one uses analysis and another agent-based simulations to provide penetrating insights.

2.2. Example 1: Monitoring and Its Efficacy

In this example, we will use the schema given above to examine the information sharing protocol recently devised by Srivatsa *et al.* [5]. For conceptual simplicity, this protocol divides time into non-overlapping *rounds* of a fixed duration T. In their protocol, each round of information exchange has four main steps: in the first step, U 'purchases' a cryptographic key $k(r_{t,U})$ where t is a round-index and $r_{t,U}$ is a measure of the self-proclaimed competency of U for not leaking information. In round t, the scheme requires that the information objects shared with U be encrypted using a parameter $R_{t,U}$ in such a way that they can decrypted using $k(r_{t,U})$ if and only if $R_{t,U} \geqslant r_{t,U}$. In the second step, U engages in information exchange and obtains information objects if $R_{t,U} \geqslant r_{t,U}$. The system limits the total number of object obtained by U in the round t to a limit $I_{t,U}$. In terms of the schema given in Figure 1, the function A is given by $A(O,U) = 1$ if $R_{t,U} \geqslant r_{t,U}$ and $\sum A(.,U) < I_{t,U}$; and 0 otherwise. We note that $R_{t,U}$ and $I_{t,U}$ are both determined in the forth step of the protocol in round $t-1$ based on monitored history of the user U.

The leakage by the user U is given by $L_{t,U} = \mathcal{B}(l_{t,U})$ where $\mathcal{B}(l_{t,U})$ is a Bernoulli random variable with probability $l_{t,U}$. In the third step, the monitoring system estimates information leakage $l_{t,U}$. The monitoring is assumed to have a delay δ_U and a random error ϵ_t, and is given by $M_{t,U} = l_{(t-\delta_U)} + \epsilon_t$. In the forth step, the protocol derives the encryption parameter $R_{t+1,U}$ and the information flow limit $I_{t+1,U}$ for the next round based on the monitored leakage, and thus completes the feedback loop. The suggested form of these parameters is given by $R_{t,U} = I_{t,U} = f(\overline{M}_{t,U})$ where f is a system-chosen linear function of past monitored leakage $\overline{M}_{t,U} = \{M_{t,U}, M_{t-1,U}, M_{t-2,U}, \ldots, \}$.

The cost charged to the user U in a round is $C_{t,U} = \lambda \cdot (1 - r_{t,U})$ where λ is a system chosen parameter. It is assumed that the leakage of a decrypted information object provides unit benefit to a malicious user while leakage of encrypted information objects provide no benefit. The net benefit is benefit minus cost. This protocol does not take into account organizational benefits or damages into account. The goal of a malicious user is to maximize its net benefit by optimally choosing the self-proclaimed competency $r_{t,U}$ and the leakage rate $l_{t,U}$. The self-proclaimed competency $r_{t,U}$ directly determines the cost of information access in each round, while $l_{t,U}$ directly determines benefits and indirectly, through the feedback of monitoring results, the total information available to it. The system can control and set the cost coefficient λ, throttle function $I_{t,U}$, and encryption parameter $R_{t,U}$. It is assumed that λ and functional dependency of $I(\cdot)$ and $R(\cdot)$ on the monitored leakage are publicly known.

In terms of the schema given above, an important, but underemphasized contribution of the work [5] is the following: by using monitored leakage as a feedback to throttle access and the rate of information access, it is possible to design an access control mechanism that reaches an equilib-

rium that does not unduely penalizes honest users who leak information inadvertently while limiting the rate of leakage from malicious users to a value that is solely a function of monitoring error. A subsequent publication [6] removes the requirement of encryption and instead replace it with a budget B that is given to each user at the beginning of a round. They study polynomial feedback and show that the leakage by malicious entities is independent of the degree of the polynomial; instead it is proportional to the product of the budget and the variance of leakage estimate. Together these studies provide an important guideline for the system design: reducing error in monitoring is more important than reducing delay in leakage estimation.

2.3. Example 2: Capping Aggregated Organizational Damage

Cheng *et al.* [4] focus on the complimentary question of how to cap the aggregated organizational damages while maximizing information flow within an organization. In the scheme proposed by Cheng *et al.*, users are issued a sum of an internal currency at the beginning of each time epoch. The scheme requires setting up a risk-token market in which the organization releases a fixed number of (say R^*) unit risk-tokens that can be traded by users amongst themselves using the internal currency issued to them. For a given access request by a user U for an object O, the access control system determines a risk value v quantified in terms of risk tokens. The user U can access O after transferring v risk tokens to the access control system. If U does not possess enough risk tokens, it may obtain additional tokens from the risk-token market by using the internal currency issued to it. Note that since a limited number of risk-tokens are released in the market, the total organizational damage is capped by R^*.

It is further assumed that information objects are used by users to generate benefits which are enumerated in terms of the internal currency. Benefits are assumed to be context dependent which evolves with time and cannot be determined *a priori*. Dynamic situation implies that the same information object becomes less or more beneficial and may cause more or less damage. Rational users try to maximize the total sum of internal currency in their possession. The approach taken by Cheng *et al.* uses agent-based simulation technique which has been applied successfully in many fields including Economics, Social Sciences, Ecology, etc.

There are two main conclusions to be drawn from this work. First, the risk-based access mechanisms are capable of capping the overall organizational damages while allowing information to be shared with users who would otherwise be untrusted with the information in a traditional access control scheme such as MLS. Second, since benefits cannot be determined *a priori*, any fixed non-transferable allocation of risk to the users will be suboptimal. However, if the users are allowed to transfer their risk budget to others who may have better opportunities to expend risk and obtain benefits, then greater overall organizational benefits can be realized.

We note that this work uses a market-based auction mechanism to determine risk allocation amongst the user population, however, its use is not mandatory to obtain benefits listed above. Market based mechanisms allocate risk with an efficiency that is close to optimal, and other mechanisms can be designed that are less efficient, but from an implementation perspective more desirable. Instead of markets, one can use Organization Theory [2] to design incentives that align individual's net benefits with that of organizational net benefits.

It is interesting to compare the approaches taken in these two complementary studies. The first work focuses on the monitoring aspect and uses analytical techniques to derive the equilibrium point of the access control system. In contrast, the second work focuses on the organizational benefits and damages, and uses agent-based simulations. Both works use microscopic models of individual behavior and provide insights into the macroscopic system design. While analysis can provide equilibrium values of the variables of interest, large scale simulations of macroscopic models of security may be the only tools available for predicting evolving complex system behaviors, convergence time, etc. Unfortunately, the acceptance of simulation based methods in the security research community remains low.

2.4. Missing Pieces

We now come back to the feasibility of the implementation of a security system that conforms to the schema shown in Figure 1. Clearly, neither the monitoring system nor the organization-wide risk-benefit tradeoffs are exploitable if the access control mechanism being used is inflexible. What is needed are systematic parameters that can be changed incrementally to fine-tune access control (see [1, 7] and references therein). The desired flexibility has many facets: (a) information sharing policies must be richer; they should base access control decisions not only on the attributes of users and information objects, but also on the attributes of the networks (social and information) within which these users and objects are embedded; (b) specifically in communication-bandwidth constrained environments, enforcement of access control needs to be decentralized yielding both energy savings and resiliency; (c) risk of information sharing must be estimated on a finer scale and it should include evolving context within which access control request is made, etc.

We believe that the Role Based Access Control (RBAC) systems lack sufficient flexibility, and the answer to more

flexible access control may lie in machine learning or data mining techniques. Conceivably,by using these techniques, systems can be adapted on-the-fly and include learning from granted exceptions, macroscopic control parameters to adjust risk calculations to reflect overall security environment, etc. A recent advance in cryptography, *predicate encryption* [3], promises to provide a crucial piece of the overall puzzle. Expressive, richer information sharing policies can be written as a predicate on user and object attributes, and information can be encrypted using predicate encryption. Later, only the users that satisfy the predicate, that is, users who are permitted by the information sharing policies, are able to decrypt and access the information object. Predicate encryption promises to provide much needed flexibility in operations since encrypted objects can encode and carry information sharing policies within themselves and do not need to be administered centrally, and thereby opening up the possibility of taking local context into consideration. We note that a time dependent parameter can be used to expire encrypted objects that have been encoded using stale policies.

Other missing pieces include implementation of risk tokens so that aggregated risk can be capped, a monitoring system to estimate information leakage, and methods to estimate risk and benefits of individual access, etc. Arguably, a catalogue of cryptographic techniques is available for implementing risk tokens and a steady progress is being made towards efficient implementation of various forms of e-currency. A monitoring system could use watermarking in conjunction with auditing to estimate information leakage. A calculation of risk and benefits is inherent in any access control system even if it is implicit and done in a static manner.

3. Conclusions

It is our hypothesis that much is to be gained by analyzing macroscopic schemata of security across system of systems. These schemata should include the concept of monitoring and its influence on user behavior, and aggregation of risk and benefit of individual accesses at the organizational level. Such schemata will allow us to study the influence of systematic controls that can be designed into security solutions to fine-tune security across a system of systems. Tools and techniques that are being used routinely and successfully in other technical fields will become available to the field of security research. To that end, we proposed a simple schema for access control systems.

To illustrate our hypothesis, we used some recent works as examples [4, 5, 6], and recast their results in a different light provided by the schema. We believe that some of the work cited here constitutes first step towards analysis of such macroscopic schemata, and overall, this is a fertile area for future research.

4. Acknowledgements

The author is Industrial Lead of the technical area "Security across System of Systems" in Network and Information Sciences International Technology Alliance (NIS-ITA). In that capacity, he benefitted tremendously from vigorous discussions amongst all security researchers participating in the Alliance from USA and UK. In particular, the author is grateful for a freely flowing exchange of ideas with Pankaj Rohatgi (IBM), Mudhakar Srivatsa (IBM), Pau-Chen Cheng (IBM), John Clark (York), Kenny Paterson (RHUL), and Greg Cirincione (ARL).

Research was sponsored by US Army Research laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] P.-C. Cheng, P. Rohatgi, C. Keser, P. A. Karger, G. M. Wagner, and A. S. Reninger. Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In *IEEE Symposium on Security and Privacy*, pages 222–230. IEEE Computer Society, 2007.

[2] R. L. Daft. *Organization theory and design*. West Pub. Co., St. Paul :, 3rd ed. edition, 1989.

[3] J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In N. P. Smart, editor, *EUROCRYPT*, volume 4965 of *Lecture Notes in Computer Science*, pages 146–162. Springer, 2008.

[4] I. Molloy, P.-C. Cheng, and P. Rohatgi. Trading in risk: Using markets to improve access control. In *NSPW*, New York, NY, USA, 2008. ACM.

[5] M. Srivatsa, S. Balfe, K. G. Paterson, and P. Rohatgi. Trust management for secure information flows. In P. Ning, P. F. Syverson, and S. Jha, editors, *ACM Conference on Computer and Communications Security*, pages 175–188. ACM, 2008.

[6] M. Srivatsa, P. Rohatgi, and S. Balfe. Securing information flows: A quantitative risk analysis approach. In *Proceedings of 27th IEEE Conference on Military Communications (MILCOM)*, 2008.

[7] L. Zhang, A. Brodsky, and S. Jajodia. Toward information sharing: Benefit and risk access control (BARAC). In *POLICY*, pages 45–53. IEEE Computer Society, 2006.