Proc. 5th RADC Workshop on Multilevel Database Security,
Fredonia, PA, October 1992

ON SOME RESEARCH ISSUES IN MULTILEVEL DATABASE SECURITY*

Ravi S. Sandhu

Center for Secure Information Systems &
Department of Information Systems and Systems Engineering
George Mason University
Fairfax, VA 22030

In this position paper I have identified some research issues which I consider to be very important for the coming decade. These are as follows.

1. Unification of integrity and confidentiality
2. Multilevel secure object-oriented database systems
3. Multilevel secure distributed DBMS architectures
4. Evaluation by parts

This list is not intended to be a complete enumeration of problems which I think are important. For instance controlling inference is clearly an important problem not included above, and there are others. Rather it is indicative of my anticipated research focus over the next 2-5 years. The rest of this paper briefly discusses each of these issues in turn.

1 Unification of integrity and confidentiality

Integrity and confidentiality have traditionally been treated as distinct objectives. In the arena of MLS Operating Systems this dichotomy has not been a problem. The integrity demands placed on an OS do not significantly conflict with multilevel confidentiality. However, for MLS DBMS's the situation is quite different. Integrity and confidentiality are often in direct conflict, as one seeks to retain the essential properties of (i) the relation model with its entity integrity and referential integrity

*A Position Paper for the Fifth Rome Laboratory Invitational Database Security Workshop, 1992.

properties, and (ii) the transaction model with its failure atomicity and serializability properties.

Polyinstantiation was proposed by SeaView researchers as one means to resolve this conflict with respect to properties of the relational model. There has been much debate about the semantics of polyinstantiation, and its necessity. I believe this debate is merely the tip of an iceberg of conflict between integrity and confidentiality. With respect to database transactions, some researchers have proposed relaxing serializability as a correctness criteria. The issue of issue of integrity constraints which cut across several security levels has been mentioned, but a careful study remains to be done.

There is a real conflict between confidentiality and integrity. In order to fully understand this conflict, and its possible resolutions, we need security models which unify these objectives in a common framework. So far, all that researchers have done is approach this issue in an ad hoc manner. At GMU we have begun a project to look at this problem in a systematic manner, with the eventual goal of producing unified models of confidentiality and integrity.

2 MLS object-oriented database systems

The need to investigate MLS object-oriented systems requires little justification. Object-oriented databases are an important technology, with considerable potential to dominate the next decade. Research at GMU on MLS object-oriented DBMS models has led to an interesting perspective on the long-standing issue of "write up." In conventional systems writing up has always presented an uncomfortable situation with respect to integrity. In relational MLS DBMS's many systems have eliminated write up due to these problems. Object-oriented systems, on the other hand, encapsulate all operations by means of message passing. Write up by means of message passing does not lead to integrity problems, since the object being written protects itself by its methods. In other words arbitrary writes are not permitted across object boundaries. So there is a comfortable feeling about write up in MLS object-oriented databases.

It turns out that to write up correctly, in terms of abstract encapsulated operations, we must execute a logically sequential computation as a collection of asynchronous computations. Now the method (call it A) which processes a write-up message may execute for an arbitrary amount of time. Moreover, the execution time can be modulated by a Trojan Horse in A. Therefore, the execution must be hidden from the method (call it B) which requests the write up (otherwise we will have a covert timing channel).

At GMU we have been looking at one solution to this problem, which is to execute A and B concurrently. The problem, stemming from integrity considerations,

is that the net effect of such an asynchronous execution must be equivalent to what is logically a sequential computation. We have shown how to ensure this in a kernelized architecture without use of trusted subjects (i.e., subjects exempted from the $\star$-property). Further study and optimization of the required synchronization protocols is in progress. We are also considering alternate solutions to this problem based on the concept of fuzzy time (introduced by DEC researchers to mitigate timing channels in an MLS Operating System).

3 MLS secure distributed DBMS architecture

The Woods Hole study identified the basic architectures for MLS DBMS's, viz., kernelized, replicated, integrity lock, and trusted subject. These concepts carry over to a distributed environment, but need further interpretation. We conjecture that the best distributed architectures will have a hybrid structure, incorporating elements of the four basic architectures mentioned above. The nature of a distributed MLS DBMS architecture will also be influenced by the nature of the network and distributed system architecture on top of which it is built. A systematic study of architectural possibilities would be a timely and valuable exercise.

4 Evaluation by parts

The authors of the TDI have consciously taken a generic (i.e., non-DBMS specific) and conservative approach to the question of evaluation by parts. This was perhaps appropriate in the first cut at this problem. However, an abstract and generic treatment leads to worst-case scenarios and pessimistic conclusions. The reality of system building and software engineering is often quite different from the hypothetical worst case scenario of the Orange Book. Trusted subjects are often used to do very specific tasks, and require very specific exemptions from the security policy of the underlying TCB. We believe that some of the pessimistic conclusions of the TDI can be relaxed in specific settings.

The authors of the TDI have also taken the approach that the Orange Book is a fixed parameter in this exercise. There has been no debate on whether or not the Orange Book could be strengthened or augmented in some way, so as to facilitate evaluation of parts. The ongoing Federal Criteria project provides an opportunity for effecting such change.

Conclusion

All four research issues we have discussed above actually transcend DBMS's. They apply to any multilevel application which has a significant integrity requirement. The DBMS arena is perhaps the one where they can be studied most fruitfully, and with the most immediate impact on implementors of multilevel applications.