REPORT ON THE 3RD NIST INVITATIONAL WORKSHOP ON INTEGRITY
Gaithersberg, MD, Sept. 26-28, 1990

Ravi Sandhu

Dr. Ravi Sandhu is an Associate Professor of Information Systems and Systems Engineering at the George Mason University, Fairfax, VA. Earlier he was an Assistant Professor of Computer and Information Science at the Ohio State University, Columbus. He has been active in computer security research beginning with his PhD thesis in Computer Science from Rutgers University in 1983. His principal research interests are: Access Control Models for Secrecy and Integrity, Multilevel Database Management Systems and Distributed Secure Systems.

This workshop was the third in a series of invitational workshops organized by NIST (National Institute of Standards and Technology) in response to the resurgence of interest in integrity following presentation of the Clark-Wilson ‘‘model’’ at the 1987 IEEE Symposium on Security and Privacy in Oakland, California. Although the stated objectives of the workshop were not fulfilled it is evident that substantial advance has been made relative to the accomplishments of the previous workshops.

The primary focus of this workshop was the creation of a document titled, ‘‘Guidelines and Recommendations on Integrity,’’ to be issued as a NIST Special Publication. NIST had earlier mailed two drafts of this document to prospective attendees. A third draft was available at the beginning of the workshop. The workshop was attended by 35-40 participants from a truly diverse collection of organizations including defense, government, commercial users, vendors and academia. There was international participation from Canada and several West European countries.

The overwhelming consensus of the attendees was that (1) there was a useful purpose to be served by such an integrity document and that NIST was on the correct track in this effort, but (2) the present draft required substantial improvement and should not be hurriedly rushed into print. The attendees expressed serious reservations about the conceptual structure of the draft. Moreover the document lacked internal consistency in its definitions and usage of terms. This alone would require a major rewrite rather than minor editing. The attendees also felt that the contents of the draft were more in the nature of a tutorial on the state-of-the-art rather than the guidelines and recommendations claimed in the title.

In spite of the criticism there was genuine appreciation of the effort NIST had expended on this task, particularly given the appalling lack of resources available to them. As a member of the planning committee I would like to specially recognize the efforts of Tim Polk and Dennis Steinauer of NIST.

Even though the workshop did not meet its stated goal of shaping the draft into a finished document, it really turned out to be a successful and productive workshop. The reason for this was the early decision by the planning committee to steer away from ''linguistic wheel spinning'' on definitions of integrity. It was instead decided to adopt a working definition for the purpose of the workshop with the acknowledgment that the working definition was not intended to be complete. This strategy neatly side-stepped the morass of issues on which the 2nd workshop of January 1989 had got bogged down.

The scope of integrity for the 3rd workshop was defined as ''ensuring that data changes only in the highly structured and controlled ways intended by the organizations.'' The draft document went on to identify three primary integrity objectives: (1) real world correspondence of data and programs, (2) ensuring that only authorized modifications take place, and (3) monitoring and accountability of authorized activities. There was surprisingly little debate about these objectives and their primary nature. The consensus achieved at this level of abstraction is by itself a creditable achievement and one that NIST should reiterate and build on in subsequent work.

The draft document proceeded in the following top-down manner. Given the three integrity objectives listed above the document attempted to identify a collection of integrity controls. Five controls were proposed as follows: (1) identification and authentication, (2) preservation of internal consistency, (3) preservation of external consistency, (4) accountability and (5) assurance. It was recognized that each integrity objective may require several of these controls. Conversely each control may support more than one objective. A particular mapping of objectives to controls was given in the draft. Finally each control was mapped to a collection of mechanisms. This final mapping was one-to-many, i.e., each control had several mechanisms associated with it while each mechanism was identified with a single control.

The quick consensus which had emerged regarding the scope of the working definition of integrity and the three primary integrity objectives did not carry over to the rest of the conceptual structure. It was generally accepted that identification and authentication and accountability are clearly identifiable integrity controls. It was also generally agreed that assurance does not belong in this list. Assurance should instead be treated as an issue orthogonal to and distinct from integrity controls. The split between preservation of internal consistency and external consistency was judged to be ambiguous and difficult to pin down. It was also considered to be of little tangible benefit to users or vendors. Many participants were specially disturbed by the attempt to assign disjoint sets of mechanisms to these two integrity controls. Several mechanisms with obvious relevance to preservation of internal consistency were listed under other categories in the draft. The participants also pointed out some assertions which were simply incorrect. The most glaring of these was the statement that ‘‘while the three requirements [i.e., confidentiality, integrity and availability] do not generally conflict, it is possible for the model used to describe them to conflict.’’ This is not true. The conflict between these three requirements is a fundamental and inherent one. The polyinstantiation phenomenon in multilevel database systems is a good illustration of this fact.

In retrospect the fundamental problem with the draft document was that it attempted to give too much structure to a body of knowledge that does not truly have that much structure at present. As a result the consensus achieved at the very high level regarding the three primary integrity objectives was rapidly dissipated as the workshop proceeded. As one group leader---Tom Chen of Wang Laboratories---put it, ‘‘Security and particularly integrity is an unbounded messy business,’’ which does not map well to an idealized vision of objectives, controls and mechanisms. This fact does confront the writers of such a document with a dilemma. Because on the other hand it is equally true that there is a bewildering variety of mechanisms which have been proposed. An exhaustive enumeration of these mechanisms without some meaningful structure serves little purpose. The point is to devise a suitable structure which has scientific and engineering merit. This is a non-trivial task which can succeed only by involving the best available talent and leading experts in this area. The process of constructing the document must allow sufficient time for debate on concepts rather than plunging into nitty-gritty details immediately.

In addition to the draft document a number of related issues were discussed at the workshop. Some of these pertained to the secondary objective of the workshop which was to address research areas and foster the development of ongoing research groups in these areas. It was acknowledged by NIST that the issues of distributed systems, integrity models, integrity metrics, database management systems and the relation between integrity and trusted systems all merit further attention and therefore fall within this category. The attendees concurred with this assessment with the exception of the relation between integrity and trusted systems. There was considerable skepticism about the merit of concepts such as ‘‘integrity covert channels.’’ Research should instead focus on the inherent conflicts between integrity and confidentiality and methods for their reconciliation.

Finally there was discussion at the workshop about ‘‘evolution of criteria’’ and how this process might develop over the next five to ten years or so. To this end two related projects were presented. The first was an NCSC sponsored study on ‘‘Integrity in the Department of Defense Computer Systems,’’ conducted by the Institute for Defense Analysis in Alexandria, Virginia. The second was an Air Force sponsored study on ‘‘Trusted Critical Computer Systems Evaluation Criteria,’’ conducted by the Information Intelligence Sciences, Inc. in Colorado. During these presentations it was evident that each of these documents was using the same words to mean different things. It was also evident that the conceptual structure used by these documents suffered from the same problems as that of the NIST draft document. There was speculation about the possibility of NIST-NCSC collaboration in developing criteria. It was suggested by some that the Orange Book might be revised and extended to incorporate integrity criteria. An alternate suggestion was to have a separate criteria from the Orange Book with separate classes and some mapping to Orange Book classes. A time frame of at least five years was suggested for this process. Others argued that there was no point in developing criteria without ‘‘worked examples.’’ Discussion on these issues was speculative at best.

At the conclusion of the workshop it was not clear what NIST intends to do next. The draft document will not be published as it stands or indeed with minor cosmetic fixes. At the same time it is not clear how NIST will salvage the worthwhile material which is in there and repackage it. Even with all its blemishes the draft document is a substantial advance over what was available at the previous two workshops. It is also encouraging to see that the Clark-Wilson ‘‘model’’ is no longer the dominant theme of this series of workshops, but that instead the focus is on basic issues with the recognition that many of the issues raised by Clark-Wilson have long been a part of standard operating practice in information systems. If NIST can invest a modest amount of resources---but greater than what it has been able to do in the past---in following up on this workshop much can be accomplished over the near term.