Proc. of the IFIP WG11.3 Workshop on Database Security,
Lake Guntersville, Alabama, Sep 12-15 1993.

On Five Definitions of Data Integrity[1]

Ravi S. Sandhu[2]
Department of Information and Software Systems Engineering
George Mason University
Fairfax, Virginia 22030-4444, USA
email: sandhu@isse.gmu.edu

Abstract This paper compares five definitions of data integrity, and shows how they can be ordered in an increasingly restrictive sequence. The most general of these, due to Courtney and Ware [6], is based on the concept of *expectation of data quality*: data has integrity to the extent that its quality meets, or exceeds, the quality requirements that users expect of it. This definition incorporates liveness requirements, whereas the others only address safety requirements. The second and third definitions are both based on the ability to modify data. One of these, due to Sandhu and Jajodia [16], defines the scope of integrity to be safeguards against the *improper modification of data*. The other narrows the scope further to safeguards against the *unauthorized modification of data*. This latter definition has been popular in recent security criteria [4,10]. The fourth definition discussed here is Biba's concept of integrity as *one-directional information flow* in a lattice [2]. We argue that Biba's definition is more restrictive than the previous three, being the one in which unauthorized modification is given a very specific meaning. The fifth and most restrictive definition comes from the network arena and requires that the *data not be modified* (or at least that any change should be detectable). The paper concludes with a discussion of how these definitions relate to other work in the integrity arena.

Keyword Codes: H.1, K.6.5, D.4.6

Keywords: integrity definitions, data quality, modification, information flow

1. Introduction

The importance of integrity as a security objective has been increasingly recognized in recent years. Yet there is no consensus on what is meant by integrity. Our objective in this paper is to reconcile five definitions of integrity that have received some prominence in the literature. Our approach is inclusive, in that we do not consider any one of these definitions to be wrong. Rather, the definitions scope the problem in different, and increasingly restrictive, ways.

[1]There was a panel discussion at the workshop based on this paper. A brief account of the panel is given in the Appendix at the end of the paper.

[2]This work is partially supported by the National Security Agency through contract MDA904-92–C-5141. The author is indebted to Howard Stainer and Mike Ware for their support and encouragement in making this work possible.

The literature recognizes a distinction between data integrity and system integrity. Data integrity is concerned with the data per se, i.e., the bits and bytes stored in the system. System integrity is a more general term, that is additionally concerned with integrity of the processing elements such as hardware, system and application software. For purpose of simplicity, our focus in this paper is on data integrity. For the most part, we assume the hardware and system software does not malfunction. System integrity will be briefly discussed towards the end of the paper.

The five definitions of data integrity discussed in this paper are as follows.

- The most general definition, which we call the *data quality definition*, is due to Courtney and Ware [6]. It is based on the concept of *expectation of data quality*: data has integrity to the extent that its quality meets, or exceeds, the quality requirements that users expect of it. The Courtney-Ware definition is the only one to incorporate liveness requirements.[3] For example, the timeliness of data may deteriorate unless the data is regularly updated. The other three definitions only address safety requirements,[4] wherein data integrity can be compromised only by an explicit act rather than by failure to act.

- The next two definitions are closely related. Both are based on the ability to modify data. One of these, due to Sandhu and Jajodia [16], defines the scope of integrity to be safeguards against the *improper modification of data*. The other narrows the scope further to safeguards against the *unauthorized modification of data*. This latter definition has been popular in recent security criteria [4,10]. We refer to these two definitions as the *data modification definitions*.

- The fourth definition discussed here is Biba's concept of integrity as *one-directional information flow* in a lattice [2]. It is more restrictive than the previous three, being the one in which unauthorized modification is given a very specific meaning. We refer to it as the *information flow definition*.

- The fifth and most restrictive definition comes from the network arena and requires that the *data not be modified* (or at least that any change should be detectable). One expects similar behavior for data on storage media. This definition is noted here for completeness, and for its position at one end of this spectrum of definitions. However, it is not discussed any further in the paper.

The rest of this paper is organized as follows. Sections 2, 3 and 4 respectively discuss the data quality, data modification and information flow definitions. Section 5 relates these definitions to other work on integrity, notably the Clark-Wilson model [5] and aspects of type enforcement [3,19]. A brief discussion of system integrity is also given. Section 6 concludes the paper.

[3]Roughly speaking, a liveness requirement says that something good *will* happen.
[4]Roughly speaking, a safety requirement says that something bad *will not* happen.

2. The Data Quality Definition

The Courtney-Ware definition of integrity was developed as a strawman for a NIST sponsored workshop on data integrity, held in January 1989 [14]. The strawman definition was published as part of the call for participation in the workshop. The strawman had been discussed amongst a smaller group of people, and represented a consensus amongst the smaller group. The objective of the NIST workshop was to develop a consensus definition of integrity. Although this objective was not met, the workshop succeeded in bringing out a number of viewpoints. One of our goals in this paper is to reconcile these viewpoints with the wisdom gained by hindsight.

Courtney and Ware [6] define integrity as follows.

> "Integrity – The property that data, an information process, computer equipment, and/or software, people, etc., or any collection of these entities, meet an *a priori* expectation of quality that is satisfactory and adequate in some specific circumstance. The attributes of quality can be general in nature and implied by the context of a discussion; or specific and in terms of some intended usage or application."

Note that this definition applies not only to data integrity, but to system integrity and beyond. As stated earlier, our focus is on the data integrity aspect.

One problem with the Courtney-Ware definition is that it (deliberately) leaves open the issue of what is data quality. It is thereby an open-ended definition. This aspect was explicitly recognized by the authors of his definition and by NIST, in that they envisaged follow-on workshops to deal with the meaning of quality in different contexts. Nevertheless, the data quality definition has a conceptual simplicity and elegance in covering a large array of concerns (beyond data integrity) in a uniform manner. Several participants (including this author [15]) agreed with this definition, although they had differing opinions about details.

The Courtney-Ware definition regards integrity as a binary attribute: either data has integrity or it does not. The NIST workshop got severely bogged on this issue. There was strong articulation of a different view of integrity as being a graded attribute, say, on a ordered scale such as very high integrity, high integrity, low integrity and no integrity. Partial orders and lattices for structured integrity measures were also proposed. Much of the motivation for a graded view stems from the influence of Biba's model [2], which will be discussed in section 4. Several participants were troubled by the lack of connection between the Courtney-Ware strawman definition and Biba's model.

In retrospect this distinction between a binary or graded view of integrity is not a fundamental one. To appreciate this consider the "thermostat model" of data integrity shown in figure 1. In this picture we show a perfect state of data integrity (i.e, one in which the data has perfect quality). The reality in any large-scale system—and these are the systems of interest—is that there will always be some deviation from this perfect state (which can exist only in some Platonic universe). Figure 1 shows the actual state of data in the system as hovering around the ideal perfect state. This figure should be visualized in three dimensions, where the actual data state never intersects the perfect ideal state

but is always some distance from it.

The binary and graded views of integrity can be explained in context of figure 1 as follows.

- *The Graded View*: In the graded view, data integrity is a measure of the deviation of the actual state of the data from the (generally unrealizable) perfect state. Continuous or discrete measures are both suitable for this purpose. Jueneman [11] presents some plausible arguments regarding how the integrity measures can be organized in lattice structure in some situations. Similar arguments can, however, be made for the applicability of almost any kind of scale in a specific situation. In other words, the graded view does not necessarily imply a lattice scale; a lattice is only one of several plausible possibilities.

- *The Binary View*: In the binary view, data has integrity if and only if its deviation from the ideal perfect state is within some a priori tolerable deviation expected by the user. Note that the measures of deviation have the same variation and range as in the graded view.

With this perspective, it is clear that there is no fundamental difference between the binary or graded views. They represent two variations on the same theme. In a sense, the graded view is more flexible, since it allows a measure of deviation from the expected quality. The binary view, on the other hand, is simpler and, perhaps, more palatable to the non-technical community. In either case, it would appear that the data quality definition has legitimate value.

Let us now consider some implications of the data quality definition. One very important consequence of this viewpoint, is that data integrity requires proactive steps to maintain data quality in addition to reactive steps. For example, say, the timeliness of data is an important quality attribute in a given context. Clearly, timeliness can be maintained only if the data is updated at some suitable rate. This requires people, or processes, within the system to undertake explicit actions. Failure to act can therefore cause loss of integrity. In general, correspondence of data to external reality requires action to maintain this correspondence whenever the external reality changes. Concerns of data quality which are internal to the database, such as entity and referential integrity in relational systems, can be maintained by reactive mechanisms. Reactive mechanisms do not initiate operations, but rather mediate the execution of operations. The other definitions of data integrity, considered in this paper, are reaction oriented. In these cases data integrity cannot be lost by itself (except due to hardware failures).

Another consequence of the data quality definition is that there should be some means of ascertaining what is a reasonable expectation of data quality. In the past, most organizations have dealt with a relatively few sources of data, about which such expectations can reasonably be built through experience. In the increasingly networked, interconnected, information-driven world of the future the sources of data may not be so familiar. In such an environment one would expect data to be tagged with some seal of approval (akin to the various quality seals one finds on consumer and business items today) indicating to the consumer some basis for assessing its integrity (be it on a binary or graded scale).

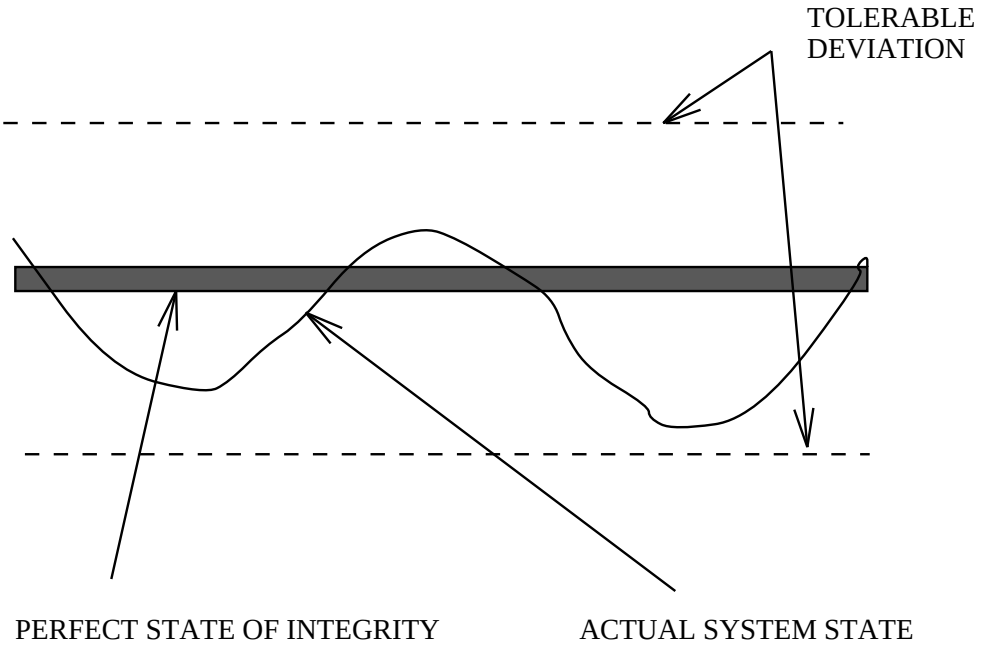

Figure 1. The "Thermostat Model" of Data Integrity

3. The Data Modification Definitions

Sandhu and Jajodia [16] propose the following definition of integrity.

> "We define integrity as being concerned with the *improper modification* of information (much as confidentiality is concerned with improper disclosure). We understand modification to include insertion of new information, deletion of existing information as well as changes to existing information."

We call attention to several important aspects of this definition.

Firstly, the definition uses the phrase "concerned with" to include both prevention and detection based mechanisms. It is evident that in distributed systems built upon insecure networks, the best one can do in many cases is to detect loss of data integrity. Even in a single computer system, improper activity of users acting within their authorizations, can only be detected rather than prevented. Also loss of data integrity due to hardware or system software failures can often only be detected (and, hopefully, corrected) rather than prevented. It is therefore important to capture both prevention (i.e., access control) and detection (i.e., auditing) mechanisms in a general definition.

Secondly, the term "improper modification" is used, rather than "unauthorized modification." This acknowledges that integrity breaches can and do occur without authorization violations, i.e., authorization is only one piece of the solution. There is ample evidence to suggest that most integrity problems arise due to well meaning actions of

authorized individuals merely trying to do their jobs. Malicious mischief by authorized individuals is also covered by the term "improper."

Finally, this definition leaves open the very important question: what do we mean by improper? The concept is that the definition of improper is really a policy decision in a given context, which intrinsically cannot have an universal answer. In this respect the Sandhu-Jajodia definition is open ended, much as the Courtney-Ware data quality definition is.

How does this definition relate to the data quality definition? The relationship comes about in the elaboration of what is meant by "improper modification." If we define improper modification to mean modification that causes data to have lesser quality than expected a priori, we have a direct correspondence between the two definitions. An important difference is that the Sandhu-Jajodia definition is only concerned with loss of quality due to explicit acts of modification; whereas the Courtney-Ware definition encompasses loss of quality due to failure to act (e.g., failure to keep the data timely and current in response to changes in the external reality being represented by the data). This is a deliberate aspect of the Sandhu-Jajodia definition. The objective is to encompass that portion of integrity which can be dealt with by access control and auditing mechanisms within the computer.

Since "data quality" and "improper modification" are both open-ended concepts, it is obvious that we can define these terms, in a given context, so that the former is always at least as general as the latter. Because the Courtney-Ware definition covers liveness requirements, whereas the Sandhu-Jajodia definition only covers safety, the former definition is more general than the latter.

The reader may have often seen definitions which are closely related to the Sandhu-Jajodia definition but use the terms "prevention" instead of "concerned with," and "unauthorized modification" instead of "improper modification." For example, the ITSEC [10] defines integrity as the "prevention of the unauthorized modification of information." The Canadian evaluation criteria similarly says, "At the present time, the Canadian definition of Integrity is that of protecting against unauthorized modification." As argued above these definitions are omitting major aspects of integrity. At the same time they are narrowing the concern to what can be accomplished by means of preventive access control. Note that these definitions are as open ended as the previous ones, in that the notion of "unauthorized" is left undefined.

The simplest definition of "authorized" is that whatever the reference monitor allows, based on (say) an access matrix, is what is authorized. This is the usual interpretation, which we assume here. One problem with the definition of "authorized" is that "improper modification" of the authorization database allows all kinds of improper authorized behavior to occur. In other words a mistake on part of the security administrator can open up avenues for improper but authorized modification of information. On this ground alone, this definition of data integrity is too narrow.

4. The Information Flow Definition

Finally, we consider a well-known definition of integrity due to Biba [2].[5] In this definition integrity is prevention of information flow from low-integrity objects to high-integrity objects. The general formulation is in context of a lattice of integrity labels, with information flow allowed in only one direction (from top to bottom). Biba's definition has had much appeal because of its close relationship to definitions of confidentiality as one-way information flow (downward) in a lattice of security labels [1,7,17].

We view Biba's definition as a particular case of the "unauthorized modification" definition of the previous section. Biba gives a particular meaning to unauthorized modification, by equating unauthorized to information flow upwards (or sideways) in the integrity lattice. As such it is much narrower than the previous definitions. One of the difficulties with Biba's definition is maintenance of an audit trail. With label-based controls the audit trail is writable (or, at least, appendable) by everybody and therefore of low integrity. This is disturbing since the whole point of an audit trail is to have high integrity. Label-based controls also do not enforce the obligation to write to the audit trail, they merely specify that it may be written.

Another aspect of Biba's model, which is often misunderstood, is its relationship to the Bell-LaPadula (or Denning) models [1,7]. In the usual formulation of the Biba model, high integrity is placed towards the top of the lattice of security labels and low integrity at the bottom. With this formulation the permitted flow of information is from top to bottom. This is in opposite direction to the permitted flow of information in the Bell-LaPadula model. Now there is nothing intrinsic about placing high integrity at the top of the lattice (or for that matter placing high confidentiality at the top). After all, top and bottom are relative terms coined for convenience, and have no absolute significance. But then information flow in the Biba model can be brought into line with the Bell-LaPadula model, by the simple expedient of saying that low integrity is at the top of the lattice and high integrity at the bottom. There is therefore no fundamental difference between the Biba and Bell-LaPadula models. Both models are concerned with information flow in a lattice of security classes, with information flow allowed only in one direction in the lattice. The Bell-LaPadula model allows information flow upward in the lattice, whereas the Biba model allows it downward. Since direction is relative, a system which can enforce one of these models can also enforce the other (given some straightforward remapping of labels to invert the dominance relation as needed).

In other words whatever data integrity can be achieved in the Biba model, can also be achieved in the Bell-LaPadula model by lattice inversion. Note that this remains true even if the Bell-LaPadula and Biba models are combined, in situations where both confidentiality and integrity are of concern. In this case we have independent confidentiality and integrity labels. However, the net effect can always be represented as one-directional information flow in a single lattice, due to the fact that the product of two lattices is itself a lattice [17].

The implications of the above discussion are that:

[5] Biba actually considers several variations of integrity. We limit our discussion to the best known of these, called *strict integrity*. The others are similar in concept.

- only a very small piece of data integrity is captured by the notion of one-directional information flow in a lattice, and

- this piece is already captured in the Bell-LaPadula model.

5. Discussion

We now discuss some integrity issues in context of the five definitions we have presented in this paper.

5.1. Other Kinds of Integrity

As mentioned earlier, there are notions of integrity that go beyond data integrity. In fact, the Courtney-Ware definition [6] enumerates some of these as follows: "data, an information process, computer equipment, and/or software, people, etc., or any collection of these entities." The terms system integrity, process integrity, processing integrity, etc., can also be found in the literature.

Our focus has been on data integrity. The data modification (and information flow) definitions could generalize to a wider context, say, system integrity, in a straightforward manner. It is not clear, however, how a modification based definition could apply to integrity of people, or even processing integrity. In short the narrow focus of the data modification definitions can be expanded somewhat, but these definitions do not have the universal applicability of the Courtney-Ware definition.

5.2. The Clark-Wilson Model

The Clark-Wilson model [5] has had considerable influence in the integrity arena. From our perspective the model deals with improper modification of data. It requires detection and prevention mechanisms, and incorporates safeguards such as separation of duties to prevent mischief by authorized users. Authorized users are also limited to use of "well-formed transformation procedure" rather than arbitrary write operations. The model does not address liveness aspects of the Courtney-Ware definition, except insofar as requiring "integrity verification procedures" to verify correspondence of data to external reality. This has the flavor of a liveness requirement, in that presumably errors detected in the stored data will be corrected.

On the whole, the Clark-Wilson model should be viewed as one approach to meeting the "improper data modification" aspects of data integrity, with a small liveness requirement attached to it.

5.3. Type Enforcement

The type enforcement model of Boebert and Kain [3,19], also captures a number of safety aspects of integrity. Type enforcement can be used to implement well-formed transformation procedures, data encapsulation, separation of duties, assured pipelines, etc., which all relate to improper modification of data. Type enforcement does not itself incorporate liveness requirements.

5.4. The Federal Criteria (Draft)

The draft Federal Criteria [9] defines integrity as follows.

> "Integrity - Correctness and appropriateness of the content and/or source of a piece of information."

Since correctness and appropriateness are not defined, this is as open ended as the Courtney-Ware definition. The Courtney-Ware definition could be argued as more general, because it is phrased in terms of data quality, which is a more general notion than the specific attributes of correctness and appropriateness.

The above formulation implies a graded view of integrity. As we have said the graded versus binary distinction is not fundamental. We feel the Courtney-Ware and draft Federal Criteria definitions are close enough that they could be reconciled fairly easily.

In the context of access control the draft Federal Criteria [9] says: "The access control objectives of organizational security policies can be divided into two classes, namely confidentiality and integrity. These objectives determine whether the organization intends to prevent unauthorized disclosure or unauthorized modification and destruction of information." We see the "unauthorized modification" definition from other criteria, as occurring here again.

6. Conclusion

In this paper we have compared five definitions of data integrity, and arranged them in an increasingly restrictive sequence. To summarize:

- the data quality definition of Courtney and Ware [6] encompasses liveness (i.e., something good will happen) and safety (i.e., something bad will not happen) requirements,
- the improper data modification definition of Sandhu and Jajodia [16] is limited to safety requirements,
- the unauthorized data modification definition popular in recent criteria [4,10] is limited to access control,
- the Biba definition [2] is limited to access control for ensuring one-way information flow in a lattice (which is not any different from lattice-based definitions of confidentiality [1,7]), and
- the network security definition which requires that the data not be modified (or at least that any change should be detectable) is the most restrictive possible.

It would be interesting to see if a similar sequence of definitions could be developed for data confidentiality.

REFERENCES

1. Bell, D.E. and LaPadula, L.J. "Secure Computer Systems: Mathematical Foundations and Model." M74-244, Mitre Corporation, Bedford, Massachusetts (1975). (Also available through National Technical Information Service, Springfield, Va., NTIS AD-771543.)
2. Biba, K.J. "Integrity Considerations for Secure Computer Systems." Mitre TR-3153, Mitre Corporation, Bedford, Massachusetts, (1977). (Also available through National Technical Information Service, Springfield, Va., NTIS AD-A039324.)
3. Boebert, W.E. and Kain, R.Y. "A Practical Alternative to Hierarchical Integrity Policies." *Proceedings of the 8th NBS-NSA National Computer Security Conference*, 18-27 (1985).
4. *The Canadian Trusted Computer Product Evaluation Criteria*, Version 3.0e, January 1993, Communications Security Establishment, Government of Canada.
5. Clark, D.D. and Wilson, D.R. "A Comparison of Commercial and Military Computer Security Policies." *Proc. IEEE Symposium on Security and Privacy*, pages 184-194 (1987).
6. Courtney, R. "Some Informal Comments About Integrity and the Integrity Workshop." *Proc. of the Invitational Workshop on Data Integrity*, (Ruthberg, Z.G. and Polk, W.T., editors), National Institute of Standards and Technology, Special Publication 500-168, September 1989, section A.1, pages 1-18.
7. Denning, D.E. "A Lattice Model of Secure Information Flow." *Communications of ACM* 19(5):236-243 (1976).
8. Department of Defense National Computer Security Center. *Department of Defense Trusted Computer Systems Evaluation Criteria*. DoD 5200.28-STD, (1985).
9. *Federal Criteria for Information Technology Security*, DRAFT, Version 1.0, Dec. 1992, NIST and NSA.
10. *Information Technology Security Evaluation Criteria (ITSEC)*, ECSC-EEC-EAEC, Brussels, June 1991.
11. Jueneman, R.R. "Integrity Controls for Military and Commercial Applications, II." *Proc. of the Invitational Workshop on Data Integrity*, (Ruthberg, Z.G. and Polk, W.T., editors), National Institute of Standards and Technology, Special Publication 500-168, September 1989, section A.5, pages 1-61.
12. Lee, T.M.P. "Using Mandatory Integrity to Enforce "Commercial" Security." *IEEE Symposium on Security and Privacy*, 140-146 (1988).
13. Roskos, J.E., Welke, S.R., Boone, J.M. and Mayfield, T. "A Taxonomy of Integrity Models, Implementations and Mechanisms." *Proc. 13th NIST-NCSC National Computer Security Conference*, Washington, D.C., October 1990, pages 526-540.
14. *Proc. of the Invitational Workshop on Data Integrity*, (Ruthberg, Z.G. and Polk, W.T., editors), National Institute of Standards and Technology, Special Publication 500-168, September 1989, section A.1, pages 1-5.
15. Sandhu, R.S. "Terminology, Criteria and System Architectures for Data Integrity." *Proc. of the Invitational Workshop on Data Integrity*, (Ruthberg, Z.G. and Polk, W.T., editors), National Institute of Standards and Technology, Special Publication 500-168, September 1989, section A.4, pages 1-14.

16. Sandhu, R.S. and Jajodia, S. "Integrity Mechanisms in Database Management Systems." *Proc. 13th NIST-NCSC National Computer Security Conference*, Washington, D.C., October 1990, pages 526-540.
17. Sandhu, R.S. "Lattice-Based Access Control Models." *IEEE Computer*, Volume 26, Number 11, November 1993, pages 9-19.
18. Schockley, W.R. "Implementing the Clark/Wilson Integrity Policy Using Current Technology," *NIST-NCSC National Computer Security Conference*, 29-37 (1988).
19. Thomsen, D.J., and Haigh, J.T. "A Comparison of Type Enforcement and Unix Setuid Implementation of Well Formed Transactions." *Proc. Sixth Annual Computer Security Applications Conference*, Tucson, Arizona, December 1990, pages 304-312.

APPENDIX

This paper served as the context for a panel discussion at the workshop. I am grateful to John Dobson, Carl Landwehr, LouAnna Notargiacomo and Marv Schaefer for their participation on the panel. I am also grateful to Carl Landwehr and Thomas Keefe for suggesting that I organize a panel session in context of this paper.

The panelists and audience by and large agreed with the ordering of these five definitions in an increasingly restrictive sequence. There were two noteworthy points which emerged from the discussion.

- It was felt that we really do not know how to apply the Courtney-Ware definition in a general setting. We have no widely used measures of data quality, and without such measures the definition is almost vacuous. In fairness to Courtney and Ware, this was recognized by them and by NIST. However, the follow-up workshops which were planned by NIST never materialized. This topic needs to be looked at by the community.

- The data quality definition impinges on confidentiality in the following way. A file labeled Unclassified but which contains Secret data clearly represents a confidentiality violation. But this can also be viewed as an integrity violation, provided we have an expectation of accurate labels. Our expectation about the Unclassified label has been violated in this example. On the other hand, if we do not expect accurate labels this will not be an integrity violation.

  This situation had generated considerable consternation among some attendees at the 1989 NIST workshop where the Courtney-Ware definition was extensively discussed. Attendees at the 1993 IFIP WG 11.3 Workshop did not find this encroachment of integrity concerns into the confidentiality domain to be troublesome, but did find it interesting and curious.