

# “User Participation”-Based Software Certification

Jeffrey Voas (jmvoas@rstcorp.com)

*Reliable Software Technologies, Sterling VA USA*

## Abstract

Except for a handful of rigorous software development schemes that are required for certification of certain critical software applications (*e.g.*, RTCA DO178-B, SAE ARP 4754, UK Defence Standard 055, and IEC 61508), there are no generally applicable software certification schemes. While there is no theoretical reason why these standards could not also be applied to less critical applications, there is a practical reason: cost. Estimates vary, but it is generally accepted that one line of safety-critical software costs between \$500-\$5,000 to design, write, test, and certify. Such levels of rigor and quality are excessive for most non-critical applications.

Nonetheless, a ubiquitously accepted process which provides guarantees about software quality for commercial software packages is still needed. These guarantees must be product-based and trustworthy. This paper proposes a certification process we believe satisfies this need. Our process collects appropriate data that can then be used to stamp limited express warranties onto mass-marketed commercial software.

**Keywords:** certification, limited express software warranties, residual testing, software users, software publishers

## 1 Introduction

Today, the number of approaches and standards for certifying software quality is increasing. These approaches either validate the integrity of the software product, development processes, or personnel [11].

The most popular approaches are process-based (*e.g.*, ISO9000 and SEI-CMM). Here, software publishers usually take oaths concerning which development standards and processes were used. Also, auditors may be employed to “spot check” a publisher’s documentation and oaths [11]. Even if a certification auditor can verify that the software publisher was truthful, that alone does not guarantee high quality software. An analogy here is that dirty water can run from clean pipes [9]. An interesting anecdote here was recently presented by Widmaier at the ISSRE’99 conference; he presented results from an experiment where a college student and professor provided a product that was significantly more reliable (using formal methods) than the same product produced by an SEI-CMM Level 4 organization (however neither product met the reliability requirements) [12].

In light of such problems, we propose a certification methodology that does not employ auditors and publisher oaths. We believe that completely *independent* product certification

is the only approach that consumers can trust. And as you will see in the next section, demands for such services are being heard from both publishers and end-users alike.

I contend that publishers should be amenable to independent agencies for this role since it allows them to avoid warranting their own software. By hiring a third party to grant software *certificates* (or *express warranties*), publishers shift this responsibility onto someone else (much like when a doctor orders a second opinion or another test). End-users will also benefit from independent agencies that offer unbiased assessments. Therefore the business case for creating independent agencies to certify software quality is strong.

We will refer to agencies that perform third-party, independent software certification as Software Certification Laboratories (SCLs). The beauty of having independent SCLs is that they provide a fair “playing field” for each publisher (assuming that each product under review receives equal treatment). A key reason, however, that SCLs have not become widespread can be attributed to the liability of being a certifier. When certified software fails in the field in a manner that was claimed to not be possible, the certifier bears some level of liability.

Why? Because SCLs represent themselves as experts. Courts in the United States are notorious for holding persons and organizations that represent themselves as professionals to unusually high standards. For example, suppose a surgeon, who never erred after 999 operations, makes a serious, *negligent* mistake during the 1,000th operation. The surgeon has a failure rate of 0.001. Even though the surgeon is highly skilled, very reliable, and has helped 999 patients, the surgeon could still be sued because of one case of negligence.

SCLs bear a similar *liability* with similar consequences for mis-certification [10]. In order to reduce an SCL’s liability, accurate methods for making certification decisions must be employed. Unfortunately, even if an SCL employs the best static analysis and dynamic testing techniques, they may still fail to consider the actual stresses that software will experience when in the hands of users. Thus SCLs suffer from the problem of accurately determining how well-behaved a software system will be in the future. And that is the key piece of information we need from certifiers.

The certification process we will propose greatly reduces this liability quagmire while also eliminating the need to dispatch human auditors. Our process harnesses the testing resources of end-users. It is similar to the approach that made Linux the most popular and reliable of all Unixes [3, 7]. And it is also similar to the Netscape Quality Feedback Agent that exists in Netscape 4.5 and the product known as PureVision. Our proposed process is totally “product-based”, i.e., it is an assessment of how well behaved the software is as opposed to how mature the development processes were that developed the code.

## 2 The Rationale Behind Product-based Certification

As software quality and information security becomes an increasing concern, the need for criteria that can accurately assess integrity and detect deficiencies grows. Voices from industry and government are starting to echo this. As an example, consider the comments of the US Department of Defense’s CIO, Mr. Money, (June 17, 1999 issue of *Federal Computer Week*):

“the quality of software we are getting today is crap. Vendors are not building

quality in. We are finding holes in it.”

Gary Beach, publisher of CIO Magazine, wrote in his April 1, 1999 column:

“Are you tired of software vendors sending you service packs to fix bugs that should have been stamped out earlier? Off-the-shelf commercial software isn’t good enough anymore. Service packs, indeed! Many CIOs I’ve talked to call them disservice packs.”

Gary Beach went on to say:

“With this column I am calling for the formation of the National Enterprise Software Certification program. As a first step I need corporate users, software vendors and value-added resellers to help me develop a rigorous certification program that would assess a participant’s knowledge of operating systems, programming languages, Internet protocols and messaging platforms. Current certification programs are too vertical. The National Enterprise Software Certification program will focus on multi-operating systems, multiprogramming languages, and multi-messaging and Internet protocol platforms. In other words, the real world of CIOs. With technology developing so rapidly, a National Enterprise Software Certificate would be good for five years. I see the program being run by a national nonprofit organization jointly funded by end users, software vendors and value-added resellers.”

The National Academy of Sciences’s 1998 “Trust in Cyberspace Report” stated:

“The absence of standard metrics and a recognized organization to conduct assessment of trustworthiness is an important contributing factor to the problem of imperfect information. In some industries, such as pharmaceuticals, regulatory mandate has resolved this problem by requiring the development and disclosure of information.”

And even Linux Inc. has felt the need to certify the interoperability of hardware and software from others as noted by the following press release:

“SAN FRANCISCO, MAY 19, 1999 Linuxcare, Inc. ([www.linuxcare.com](http://www.linuxcare.com)), the leader in technical support, consulting, education, and product certification for Linux, announced today the launch of Linuxcare Labs, a comprehensive program for independently certifying hardware and software for use with the Linux operating system.”

While these comments reveal a heightened dissatisfaction with existing quality, they only represent one reason why I believe that product-based certification is inevitable. The other reason stems from legislation (*e.g.*, the Y2K liability limiting legislation and the Uniform Computer Information Transactions Act (UCITA) that was ratified by the National Conference of Commissioners on Uniform State Laws (NCCUSL) on July 29, 1999). Without going into UCITA’s details, consider excerpts from Barbara Simons’s (President of the ACM) July 12, 1999 letter to the NCCUSL:

“The computing professional must strive to achieve quality and to be cognizant of the serious negative consequences that may result from poor quality in a system.” The UCITA will hinder this pursuit, and result in possible harm to the general public. UCITA makes it too easy for software publishers to avoid facing any legal consequences for defective software.”

While UCITA is still legislation and not yet law, if it were to become law in each of the 50 states, software publishers will enjoy even less accountability than they do today [8]. The question then becomes: “where will consumers turn if they are unable to accept low quality commercial software?” I believe it will be to organizations that can offer assurances that are independent from those claimed by the software publisher.

Thus I believe that the stage has been set for independent “product-based” software certification. The question, however, is in what form will such services be offered? And what assessment technologies should be used? A single certification methodology is probably not going to “fit all” types of software. Numerous certification methodologies will be needed and each must be designed according to: (1) what the software does, and (2) what level of integrity is guaranteed by the certificate’s “seal of approval.”

### 3 Our “Limited” Software Warranty Model

Our interest is in certifying software applications and components that are mass-marketed and have the potential for many different types of usage. Our short-term interest is in certifying applications such as Microsoft Word, Powerpoint, Excel, etc. Our long-term interest is in certifying smaller components that could be embedded in a variety of applications. Our goal is to provide a certification process that fosters greater software reuse for different types of software.

Component-based software engineering (CBSE) is simply building software systems from software parts. But the virtues of CBSE have been touted for years. And there is *little* evidence that we are close to CBSE becoming the de facto standard for how software systems are built. Why is this when glue technologies (e.g., CORBA, ActiveX, COM, DCOM) for component interoperability and reusable component libraries already exist?

In our opinion, CBSE has not become ubiquitous because of: (1) widespread distrust of software components, and (2) a lack of knowledge about how to design-for-reuse. Any two components can be glued together rapidly, but that is not sufficient for popularizing CBSE. Integrators must have confidence that the right components were glued, they were glued together correctly, and the components are dependable.

And as already mentioned, trust in a component’s dependability must come from someone other than the software publisher. We have already listed two reasons for this: the liability associated with making claims concerning quality, and end-user distrust of publishers. Two additional reasons for not allowing the software publisher to be the person to assess the goodness of their products are: (1) software publishers are unlikely to have adequate resources to test to the levels that would justify software warranties, and (2) publishers are likely to make incorrect assumptions about how users will use the software.

Interestingly enough, SCLs that might hope to employ in-house testing as a means to certify software will suffer the same problems just mentioned: inadequate levels of testing

and incorrect assumptions. In fact, the degree of testing that can be performed in a cost-effective manner at a SCL may be less than that already done in-house by the software publisher. Thus the idea then that SCLs can do in-house testing as a means for granting warranties is dubious.

In our opinion, software warranties must be the end-result of massive amounts of operational usage. Such testing will have demonstrated product stability in fixed environments and fixed marketplace sectors (embedded, desktop, Web-based, etc.). Without this, these limited warranties will either be too restrictive or unbelievable.

If SCLs and publishers cannot perform adequate product testing, who can? Our answer is the disorganized body of software testers known as the users. They can be unified to overcome this problem and do so in a way that is advantageous to them. The issue then is how to best tap into them in a manner that makes software warranties into a reality.

We propose user-based product certification as a means for granting software warranties. Our approach collects valuable field data in a non-intrusive manner from the user's environment with their permission. Currently such data is rarely collected.

Our approach will exploit testing technologies similar to the *residual testing technologies* discussed by Pavlopoulou [6]. Residual testing is post-release testing that tests fielded software as it operates. Residual testing employs instrumentation that is embedded in operational software to collect information about how the software behaves. This information can reveal what software code is executed, whether assertions fail, and whether the software itself fails (this can be accomplished by bundling in a limited regression testing capability). We will also collect information on how the product is used. We can measure the frequency with which certain features are called, simply collect files of inputs (that later can be used to build operational profiles), etc. These sets of information will serve as the basis for issuing *limited* software warranties.

Our proposed certification process is shown in Figure 1. Here, a software publisher subjects their finished product (i.e., the post- $\beta$  version) to a residual testing tool which creates one instrumented copy (non-identical copies could also be built but that will make the tasks of the SCL more complicated). The copies are then supplied to the SCL. The SCL provides the instrumented versions to pre-qualified users from different marketplace sectors. Pre-qualified users will be those who will use the product in a manner consistent with how the SCL wants the product used. (Microsoft's beta-user program is an excellent example of how to screen possible users to those who will actually provide the greatest amount of information.) Periodically, the SCL gathers the information from the user sites.

All user information is merged and statistics are computed by the SCL. The SCL provides: (1) statistical data back to publishers concerning how their product was used, and (2) data telling publishers how their products behaved in the field (i.e., product quality). Statistics must be generated such that a backwards trace to any specific user is impossible. This is of paramount interest. (We will later look at why previous attempts by publishers to force users to reveal similar information were unsuccessful and how our model avoids those pitfalls.)

Recall that the goal of this process is for an SCL to confidently determine whether a product should be warrantied and what it should be warrantied for. That is, which platforms, which market sectors, and which environments? This is why we refer to our warranty as "limited." As an SCL gathers additional data over time, the SCL may be in a position to broaden a warranty, i.e., make it into a fuller, less restrictive warranty.

### Figure 1: **Basic Certification Process**

The motivation for our approach to software warranties stems from a client request some years back. They wished to employ the AIX operating system in a highly safety-critical environment. The client needed to know how reliable different AIX operating system and library calls were. The client could not find the information (nor did IBM know) even though thousands of AIX users worldwide had generated enough anecdotal evidence that when taken together, almost certainly would have answered our client's question. But because that information was lost, the client had to spend substantial resources re-testing each operating system and library call. And worse, it is possible that other organizations are redoing what our client did. This is clearly wasteful.

Note that we are not suggesting that you seek volunteers to fly uncertified software-controlled aircraft or use uncertified software-controlled medical devices in order to see if the software in them should be certified. Instead, what should occur is for the software to be first certified for non-critical market sectors. Then, if a supplier of a safety-critical product were to discover that this certified software utility was certified for an environment that mirrors the safety-critical product's environment, then the supplier has the option to employ it. (Thus all we have said here is that the operational profiles are equivalent between the non-critical and safety-critical environments.) If the software behaves well in this and other safety-critical applications, and if enough of this evidence can be collected demonstrating confidence in the software utility, then evidence will exist that warrants certifying the utility for the safety-critical market sector. Clearly much caution must be taken here, but recall that we are dealing with post- $\beta$  software to begin with. We are never dealing with untested software, even when a piece of software is first embedded into a safety-critical environment.

## 4 Key Characteristics of Our Model

We have described our model. For our model to be acceptable to users, SCLs, and publishers, we deliberately built in the following characteristics:

1. Certification is performed on a stable, non- $\beta$  version of the product.
2. Certification is not done by the software publisher. The publisher delivers a fixed number of versions of a product to the SCL. These versions have residual testing capabilities built-in. The SCL licenses the versions to pre-qualified users from various application domains (e.g., safety-critical, desktop, Web, etc.).
3. The testing methods and type of usage information collected must be determined by the SCL. This information should be encrypted during transfer back to the SCL and details concerning the type of information being collected must be made available to participating users in order to diminish concerns regarding user privacy.
4. Certification decisions are based on significant operational experience created by the user base. This approach employs orders of magnitude more testing than the publisher or SCL could perform. Also, certification decisions are based on real operational scenarios as opposed to hypothesized scenarios. Hypothesized scenarios are usually the best that a SCL and publisher can base pre-release testing on. That will rarely be good enough, for example, to convince an actuary in an insurance company to insure a software product. Because our approach employs field data, the data an SCL will collect could be used by insurers that hope to insure software-based systems. Also, insurers may feel confident enough to indemnify an SCL against mis-certification since each warranty is based on historical data.

We should also mention that certification decisions may also need to be based on the environment and configuration under which the software is operating. This adds one additional set of information that the user may need to supply to the SCL or that the instrumentation may need to supply when results are passed back.

5. User data is collected by background processes that minimally interfere with users. Users expend no resources to trap and collect the information (other than the computer resources necessary to gather and store the information). Although users are not required to perform manual tasks to collect the information, their versions will have performance degradation due to the background processing cycles required to perform the post-deployment testing. Users that participate in this accelerated testing program will be compensated with reduced rate or free software.
6. All testing data collected from users is delivered exclusively to the SCL. The raw information does not go to the publisher. SCLs agree to legal confidentiality between themselves and users. Once the raw information is scrubbed and user identities are not traceable from the composite information, the composite information is passed back to publishers. This allows them the opportunity to improve quality on future upgrades and focus in-house testing toward their user base.

7. A software warranty is limited to different market sectors (safety-critical, desktop, e-commerce) and different environments (Unix, Windows-NT, Macintosh). When enough data is collected from the field affirming that a component is working properly in a particular market sector, the SCL will provide software warranties specific to that market sector. For example, a SCL warranty might read as: software product  $X$  is warrantied to perform with a reliability of 99.9 in the Windows-NT desktop environment.

What constitutes “enough data” is a function of basic statistics and the type of certification being sought. For example, to be 99% confident that the true probability of failure of a system is less than 0.00001, 460,000 test cases are needed. And so we will rely on basic mathematical principles before assigning warranties.

8. Users *consent* to participate. Users that do not wish to participate in this process do not have to. They are free to license non- $\beta$  versions that do not undergo residual testing.
9. Publishers compensate SCLs for their services. The SCL will also own the collected data. Because the data has economic value, fees placed on publishers by the SCL should be reasonable. Note, however, that publishers are gaining invaluable data on how the software is being used and gaining enormous amounts of product testing. Therefore it is reasonable that they pay a fair price.

We believe that there is a reasonable business case for the SCL selling the data in a format like that found in *Consumer Reports* magazine). Recall our AIX story where a client had to re-test the software to get the data it needed on AIX’s dependability, and other organizations will have to do the same. Why not create the information once and then sell it? Also, it is plausible that a non-profit company or government agency serve as an SCL.

10. No auditors are needed. Most existing certification models require that trained auditors visit publisher sites and sift through requirements documents, test plans, etc. Our scheme avoids the pitfalls associated with human auditor error.
11. Unlike ISO-9000 and CMM, our product-based certification allows innovative processes to emerge. If a developer organization can produce certifiably good products that conform to an in-house standard, then they have developed new processes that other development organizations should consider. After all, ISO and CMM may lock us into old process ideas prematurely.
12. This certification process creates a framework for performing *accelerated testing*. Accelerated testing requires the distributed environment defined in this model. Note that accelerated testing is frequently necessary for the highest reliability goals (as defined by Butler and Finelli in [4]). And it can benefit reliability models such as Voas and Miller’s Squeeze Play; that model often requires massive amounts of testing to overcome low testability scores [5].
13. This certification process provides another benefit to publishers: fewer “point” releases. Once their product has undergone certification and succeeded, there should be fewer releases of versions for the purpose of “bug fixes”.

14. If software packages can be confidently warrantied for certain environments and under specific assumptions, then users with equivalent environments no longer to invest heavily in re-testing the licensed software.
15. Because these methods are automated, this certification approach is *repeatable* and *reproducible* (meaning that different SCLs will produce the same warranty because they would receive the same fielddata).
16. And finally, this process forces publishers to define what is correct behavior for their software. This alone could result in higher quality software products.

Clearly there are benefits to our approach. But those benefits cannot be realized until we know precisely what data is sufficient to justify issuing limited software warranties. We envision it including reliability assessments, data from assertions, monitoring exception calls, data on operational usage, and code coverage analysis. But ultimately it depends on how strong of a warranty is sought.

Other analyses are possible, but these seem to be the more important indicators of whether a component has been exercised thoroughly enough to warranty it for specific environments and market sectors. This is a topic for future research. Our immediate plans are to work with technology insurers to select the residual testing technologies that provide data sufficient for offering software insurance premiums and software warranties.

One other issue that we should mention is how will we know that the instrumentation added by the tool is sufficient. In short, how do we certify that the tool is adding the correct instrumentation? To address this, we consider this problem no more difficult than how other software tools are qualified in various “regulated” domains, like commercial avionics. For example, the FAA has an official procedure for software tool qualification. A similar scheme can be applied here to ensure that the tool provides correct instrumentation. And it will be the responsibility of the SCL to ensure that the user of the tool (*e.g.*, the publisher) adds adequate instrumentation. Unscrupulous software publishers might attempt to force the tool to add “bogus” instrumentation that fails to perform testing in order to gain quick approval. It will be the responsibility of the SCL to mitigate this possibility.

## 5 Related Approaches

While our proposed model for combining SCLs with residual testing technologies is unique, it admittedly leverages existing and previous ideas. More specifically, it is conceptually quite similar to Netscape 4.5’s Quality Feedback Agent and the product known as PureVision. In this section, we will discuss these and other related ideas.

Our first example is from a product called PureVision that was released in 1995. PureVision was offered by Pure Software. The product performed crude residual testing functions. It worked in a manner similar to how we have defined residual testing: a publisher produced copies of a product that were able to monitor themselves at user sites [1]. The copies sent back a report *to the publisher* concerning: (1) which user and user site were using the product, (2) the version number, (3) system configuration, (4) when the version started executing and stopped, (5) which program features were used, and (6) the amount of memory used

at exit. If the product failed, exit codes and a stack dump were added to the report. Pure knew that users would be wary of publishers looking over their shoulders and thus included an option whereby a user could inspect a report before it was sent back. An option to not send a report back was also available. According to former Pure employees, speculation for why PureVision did not survive is that users were unwilling to provide such detailed, non-technical information (e.g., which user was using it, at what times, on which host, etc.). In contrast, the information we will collect is more technical; our approach spies on the software and its correct behavior as opposed to spying on the user.

Our second example is similar to PureVision. Netscape 4.5 contains an option called The Netscape Quality Feedback Agent. The agent sends feedback to Netscape's developers concerning how the product is performing. The agent is enabled by default and is activated when Communicator encounters some type of run-time problem. When the agent is activated, it collects relevant technical data and displays a form in which a user can type comments. Netscape uses this data to debug known problems and identify new ones. Unfortunately, however, most users do not activate this feature for similar reasons to those for why PureVision did not survive.

Our next example is taken from the Software Testing Assurance Corporation (STAC) which was a for-profit venture started in 1997. STAC attempted to be the sole certifier for Year 2000 insurance on behalf of insurers. STAC's certification procedure was based on a public-domain testing standard (for Year 2000 remediated code) and STAC-approved auditors who were to make site visits to companies seeking Year 2000 software insurance. In the end, however, these insurance policies were never offered to my knowledge.

Microsoft has long employed beta-testing as a way to collect information on how their products perform in the hands of users. And for this information Microsoft compensates users. Pre-qualified users are given advanced copies of a product at reduced rates in exchange for feedback concerning product stability, usability, and reliability. Microsoft uses this information to decide when a product is ready for general release. This demonstrates user willingness to participate in exchange for discounted software. Also, Cusamano and Selby have claimed that Microsoft gathers detailed user profiles from specifically instrumented versions of their software [2].

And finally, consider Linux, a Unix operating system project that began in 1991 and today has nearly 8 million users. Linux is the product of hundreds of users, all of whom donated their time to write the system. Their efforts paid off: Linux is considered to be the most reliable of all Unix operating systems [7]. In fact, the success of Linux is often used as the argument for why products with open source are the only products that are trustworthy.

So what can we surmise from these anecdotes? From Microsoft and Linux, we see that users are willing to participate in efforts that result in improved software. From PureVision and Netscape, we see that publishers have a serious desire to attain field information. And from STAC, we see interest from industry in forming for-profit software certification corporations.

What still remains to be seen, however, is whether users will trust an SCL to act as an intermediary between themselves and a software publisher. As long as the confidentiality agreements between: (1) the SCL and publisher, and (2) SCL and user are binding, we believe users will.

## 6 Summary

Demands for SCL services translate into business opportunities for those that can overcome the liability risks. In fact, several SCLs are already in existence today. For example, KeyLabs certifies language purity for Java but makes no promise about software quality. Because language purity is trivial to test for, the liability for KeyLabs is low.

Our approach dissolves the distrust associated with schemes that use auditors or measure process maturity. We capture user information that is traditionally discarded: (1) is the software behaving correctly?, and (2) how is the field using the product? Our approach does not invade user privacy; participant users agree to license instrumented versions from intermediaries who sit between them and the publishers and participant users are informed as to what information is collected.

We envision that our proposed approach will initially be applied to complete applications (*e.g.*, Microsoft Word). Once it has been applied at that level of granularity and we have a better idea as to how to implement this scheme, it should be ready for smaller sized objects (*e.g.*, plug-ins and components). This clearly has the potential to decrease the cost of software development by fostering CBSE.

And finally, the time it takes to attain a limited software warranty can be reduced by attracting more participant users and by only releasing instrumented versions to persons that are committed to seriously stressing a product. This will reduce the time until all users can enjoy un-instrumented, warrantied software.

## Acknowledgements

I'd like to thank the anonymous referees for their comments on this article. I'd also like to thank Chuck Howell, John McManus, Roger Alexander, Dolores Wallace, Frank Ackerman, Lora Kassab, and Keith Miller for their comments on earlier versions of this article.

## References

- [1] L. BINGLEY. PureVision: Shedding Light on Black Art Betas, APT Data Services, Inc, New York, June 1995, page 17.
- [2] C. CUSAMANO AND R. SELBY. *Microsoft Secrets*. The Free Press, 1995.
- [3] B.P. MILLER ET AL. Fuzz revisited: A re-examination of the reliability of UNIX utilities and services. Technical report, University of Wisconsin, Computer Sciences Dept, November 1995.
- [4] R. BUTLER AND G. FINELLI. The infeasibility of experimental quantification of life-critical software reliability. In *Proceedings of SIGSOFT '91: Software for Critical Systems*, pages 66–76, New Orleans, LA, December 1991.
- [5] J. VOAS AND K. MILLER. Software Testability: The New Verification. *IEEE Software*, 12(3):17–28, May 1995.

- [6] C. PAVLOPOULOU. Residual Coverage Monitoring of Java Programs, Master's Thesis, Purdue University, August, 1997.
- [7] B. P. MILLER, L. FREDRIKSON, AND B. SO. An empirical study of the reliability of unix utilities. *Communications of the ACM*, 33(12):32–44, December 1990.
- [8] C. VOSSLER AND J. VOAS. *Advances in Computers*, chapter Defective Software: An Overview of Legal Remedies and Technical Measures Available to Consumers. Academic Press, to appear in 2000.
- [9] J. VOAS. Can Clean Pipes Produce Dirty Water? *IEEE Software*, 14(4):93–95, July 1997.
- [10] J. VOAS. Software Certification Laboratories: To Be or Not to Be Liable? *Crosstalk*, 11(4):21–23, April 1998.
- [11] J. VOAS. The Software Quality Certification Triangle. *Crosstalk*, 11(11):12–14, November 1998.
- [12] J. WIDMAIER. Building More Reliable Software: Traditional Software Engineering or Formal Methods? In *Proceedings of ISSRE'99 Industry Practices and Fast Abstracts*, pages 253–260, November 1999.