

Computer-Related Risk Futures

Peter G. Neumann*

Abstract: This paper reflects on many risks in the development and use of computer-related systems. It considers past and future alternatives, suggests some remedial approaches, and offers a few broad conclusions. Various long-touted common-sense approaches that are holistic and proactive are more urgently needed now than ever before.

Keywords: Risks, trustworthiness, security

1 Introduction

The word *futures* in the title of this paper is intentionally ambiguous. On one hand, we might imagine many strategies for alternative visions of future times in which risks could be controlled, managed, dramatically reduced, or even in some cases effectively eliminated, perhaps by various means considered here. On the other hand, *futures* suggests some reactive business-as-usual interpretations. For example, you might think of an exchange marketplace that could trade in expectations relating to risks—perhaps using actuarial tables on past and expected losses to issue, buy, and sell insurance policies against computer-related disasters, or speculating by short-selling stocks in anticipation of (or prior to intentionally causing) malware exploitations and other serious cyberattacks. Further derivative activities might also come to mind, such as trading within caps on liabilities, somewhat akin to using carbon-cap trading in an attempt to address climate control.

Oversimplifying somewhat, proactive strategies should constructively appeal to enlightened people in R&D and operational communities. Reactive strategies might rather destructively appeal primarily to speculators in financial markets and short-term profit seekers. Of course, the former approach is recommended here, and the latter approach is mentioned only as a dead-horse straw herring that needs to be whipped in the bud (and nipped in the budget).

*Computer Science Laboratory, SRI International, Menlo Park CA 94025-3493, Neumann@CSL.sri.com

This author [8, 10] and colleagues (such as Bob Boyer, Fernando Corbató, Peter Denning, Edsger Dijkstra, Virgil Gligor, Tony Hoare, Nancy Leveson, J Moore, David Parnas, John Rushby, and Jerry Saltzer, to name just a few) have long advocated principled approaches whose adoption could seriously reduce many of the risks. Unfortunately, during roughly the past 40 years much of that shared wisdom (e.g., [4] from 1968 and [8] from 1969) has been largely ignored in practice. This wisdom certainly bears respecting in the persistent hope that it might eventually gain greater traction, particularly in developing greater awareness of the risks of inaction.

2 Historical Perspective

Many application areas such as critical infrastructures, real-time control, and computer-aided elections require trustworthy computer systems satisfying requirements such as predictable security and reliability, as well as enterprise survivability and manageability in the face of a wide range of threats—including hardware malfunctions, penetrations, denial-of-service attacks, insider misuse, malware, and improper human behavior. However, many existing systems are inadequately trustworthy for such applications. Further, overcoming a pervasive absence of meaningfully trustworthy systems is only part of the problem, because many of the problems have serious nontechnological aspects. Overall, the risks continue to be problematic, with rampant opportunities for harmful consequences—as the variety of uses expands, demands for automation increase, technological complexity grows, and capabilities of would-be adversaries outpace those of defenders.

Characteristically, risks continue to recur that violate trustworthiness requirements such as those noted above. They tend to repeat themselves in one form or another, suggesting a lack of foresight. For example, see 33 years of ACM SIGSOFT's *Software Engineering Notes* (www.sigsoft.org/SEN/), 24-plus years of the ACM Risks Forum (see www.risks.org)

for archives and how to subscribe), and Computer-Related Risks [9], for copious examples, along with an annotated index to many past cases [15]. The series of *Inside Risks* columns in the *Communications of the ACM* succinctly summarizes many of these risks (www.csl.sri.com/neumann/insiderisks.html). A sample set of problems from that source material appears in my 2006 ACSAC paper [12] on the risks of untrustworthiness, which considers just a few prototypical problem areas—notably, unreliable backup, nonrobust networking, and systems that are unsafe, unsecure, or both.

Reflection on the RISKS archives leads to several observations. The very first issue (volume 1, number 1) on August 1, 1985 noted that hot topics included the relationship of computer systems to money, privacy, and elections, and anticipated risks in defense systems, human safety, consumer protection, and health care. That first issue had items on the Strategic Defense Initiative—the anti-missile missile system euphemistically known as Starwars—including a summary of Dave Parnas’s analysis of why that approach was very unlikely to succeed, and an item on Dave’s resignation from the advisory panel. It also noted prior events including two serious automobile recalls due to program bugs (the El Dorado brake computer and the Mark VII computerized air suspension) and at least two heart pacemaker problems (one of which resulted in a death).

In crafting some expectations for the future in that first issue, I reported on a variety of money losses, security problems, and a string of potential misuses and undetectable mishaps in elections that suggested opportunities for Trojan Horses and fraud. On the subject of elections, I cited an article by David Burnham that appeared on July 29, 1985 in *The New York Times*, on vulnerabilities in various computerized voting systems. Burnham noted that about 60% of the votes were then being counted by computer programs, with over a third of those votes being counted by one program (or variants thereof) written by what was then Computer Election Systems of Berkeley CA (subsequently merged into Business Records Corp., with legacy ties to today’s ES&S—which is now seeking to acquire Diebold’s voting subsidiary). That system reportedly could be undetectably manipulated to modify election results. Burnham wrote, “The allegations that vote tallies calculated with [that program] may have been secretly altered have raised concern among election officials and computer experts. ... In Indiana and West Virginia, the company has been

accused of helping to rig elections.” At the time, I wrote, “This topic is just warming up.” Twenty-four years later, it remains very hot.

For two historically separated datapoints on the history of that topic, see Ronnie Dugger’s article in the November 7, 1988 issue of *The New Yorker* and his later reprise in the August 16/23, 2004 issue of *The Nation*. The lack of fundamental progress reflected during the 16-year interval between those two articles (both of which are on my website) and the continuing lack of integrity in the use of all-electronic (paperless) voting systems are very troubling.

Volume 1 of RISKS is replete with thoughtful items with considerable foresight on risks that needed to be addressed. Later volumes considered more of the same risks, sometimes in new manifestations—and particularly problems relating to security and privacy, reliability and safety (for example, in transportation, medical applications, and control systems), and human well-being. Further, numerous cases have been recorded in the RISKS archives that involve extremely bad system development and software engineering practice, poorly designed human interfaces, human errors, and malicious misuse by both insiders and outsiders. Looking back, very few topics of concern today were unanticipated, at least conceptually.

Not surprisingly, all of the above topics are still being discussed in RISKS, including the continued and much-spirited controversy over election systems. New incarnations of anti-missile missile systems keep appearing, although the technology and software engineering assumed to someday exist still seem far away. Widely propagating electrical power outages have occurred surprisingly frequently since 1965. Safety risks in aviation, railroads, process control systems, nuclear power, health care, and many other application areas are threaded through the entire RISKS archives, along with myriad security and privacy problems. The 1980s saw considerable attention devoted to multilevel-secure systems, which still seem like a pipe-dream with respect to the likelihood of high-assurance commercially available products. The 1990s and 2000s were marked by beliefs in some circles that off-the-shelf products would eventually be sufficiently trustworthy, with considerable empirical evidence to the contrary along the way.

What seems to have changed somewhat over the years? Local security and privacy problems are increasingly arising more globally in new contexts, such as health care, distributed systems, Web servers and browsers, cloud computing, and Internet-connected,

previously isolated systems as in national infrastructures that are inadequately prepared to meet the increased risks. We have widespread cases of malware, identity fraud, spam, phishing attacks, and distributed denial-of-service attacks. There is also increasing awareness that cryptographic algorithms are by themselves inadequate to solve security problems, and ultimately must rely on being embedded in trustworthy hardware and software.

Various common themes have emerged over the years. For example, many usability problems have resulted from poorly conceived human interfaces that place inordinate demands on timely human intervention. However, blame for serious risks is typically placed on system operators, pilots, and users, not on the designers and implementers of systems with inadequately specified requirements, badly conceived architectures and human interfaces, poor development methodologies, sloppy software engineering, inappropriate programming languages, and so on.

Further, short-sighted local optimization often results in risks that could have been avoided with even a smidgen of long-term planning and globally motivated considerations. Badly chosen tradeoffs may actually result in reducing trustworthiness. Many issues are all too often ignored, such as long-term enterprise survival and privacy.

Many of the risks discussed in my 1995 book (Computer-Related Risks [9]) are still recurring in one guise or another, and many of its recommendations remain relevant but unheeded. We continually tilt at the same windmills, and the windmills seem to be winning. It is clear that some far-reaching proactive measures are urgently needed—although this is not a new conclusion. Indeed, we are in some ways merely reiterating certain previous findings and recommendations, such as those found in reports of the National Research Council relating to trustworthiness—from 1983 (Multilevel Data Management Security [21]), 1990 (Computers at Risk [3]), 1998 (Trust in Cyberspace [22]), to 2007 (Toward a Safer and More Secure Cyberspace [5]).

3 Some Pressing Problems

Consider just a few representative pandemic problem areas that encompass a variety of application scopes, complexities, and potential risks.

- **Critical Infrastructures.** The November 1965 Northeast power blackout has been followed by nu-

merous similar events. Testimony before the Clinton-era President's Commission on Critical Infrastructure Protection repeatedly suggested that all of the critical infrastructures were vulnerable to attacks on the survivability, integrity, and security of their information infrastructures. In 2009, vulnerabilities in power distribution and other national critical infrastructures have been reported in public forums as being widespread and relatively easy to exploit, and computer systems controlling power grids have been victimized by inserted trap doors and Trojan horses. Further, the ability to detect, diagnose, and respond to attacks and outages is inadequate. The risks abound.

- **Malware.** Perhaps the most pervasive risks involve the continued lack of trustworthiness within critical information system and network infrastructures. For example, various generations of the Conficker malware suite [17] emerged during 2009, initially exploiting unpatched operating systems. Conficker represents some of the most sophisticated malware seen thus far, morphing in stages into successively more devious functionality. And yet, the potential for much more serious consequences is almost untapped, including financial ruin, corporate demise, clandestine surveillance, phishing attacks, massive identity fraud, and so on. This is a problem that has long been recognized and typically neglected. For example, the 1988 Internet Worm provided a harbinger of possible malware that has progressively led to viruses, worms, and other malicious code. But flaws such as buffer overflows, overly permissive `.rhosts` files, and easily compromised reusable passwords are still rampant two decades later.

- **Electronic Voting.** Although elections are one of our most critical infrastructures in maintaining democracy, irregularities in elections are perhaps as old as elections themselves. Serious vulnerabilities exist in election processes today, and risks in the information system infrastructures are only one area of problems. The frenzy after the 2000 U.S. national election resulted in an aggressive move toward all-electronic systems. However, all proprietary paperless electronic systems that emerged have pervasive security vulnerabilities (e.g., see the California Secretary of State's 2007 Top-to-Bottom Review, www.sos.ca.gov/elections/elections_vsr.html), almost no meaningful accountability, and inadequate oversight and governance. These systems have seriously exacerbated the overall lack of election integrity, accompanied by inequities in voter registra-

tion, sloppy management of registration databases, and politicization of election processes. A recent example involves the 2009 indictments of five people in Kentucky for conspiracy to commit vote fraud, extortion, and tampering with grand jury witnesses in a subsequent attempt at a cover-up during elections in 2002, 2004, and 2006. Insiders exploited a misleading user interface to alter votes before ballots were actually cast. Indeed, today's electronic voting systems are held to much weaker standards than (for example) gambling machines and lotteries. Worse yet, Internet voting is a potential disaster in waiting; although highly desirable in principle, it requires much more selective use (perhaps only for acquiring an appropriate ballot, but not for casting it?), security controls, and oversight before it might be considered trustworthy. Approaches to election integrity are needed that span the entire process from beginning to end, including the entire supply chain. (For example, see [7], which represents requirements for election integrity within the Common Criteria framework.)

4 Would-be Remedies

Various approaches to risks have been postulated. A quasi-analytic approach involves risk assessment that identifies and attempts to quantify the most serious risks, sometimes based on false or incomplete assumptions. However, such analyses are typically followed by so-called risk management that often ignores the results, declares them sufficiently unlikely, seeks insurance coverage, or perhaps attempts to reduce or prevent the causes of those potential risks. Another more far-sighted approach involves constructive and proactive action as suggested above—for example, developing systems designed to satisfy well-specified and realistic requirements, with soundly layered and composable system architectures (see fm.csl.sri.com/LAW09/#program for recent papers on the 2009 workshop on this topic, including [14], and John Rushby on what we can learn from other disciplines), principled software engineering practices, inherently safer programming languages, and intelligent system administration.

A document developed for Doug Maughan at the Department of Homeland Security, A Roadmap for Cybersecurity Research [6], considers 11 areas of hard problems in information security: scalable trustworthy systems, enterprise-level metrics, lifecycle of system evaluation, coping with insider threats and mal-

ware, global-scale identity management, survivability of time-critical systems, situational awareness and attack attribution, provenance, privacy-aware security, and usable security. Each area needs significant effort toward research, development, evaluation, and technology transfer. Of considerable importance is the extent to which these 11 areas are interrelated, which strongly suggests the need for holistic approaches (e.g., [11]). Although decoupling these areas may be desirable for many reasons, the interdependencies must nevertheless be realistically accommodated throughout R&D, including the establishment of requirements, architectural design, implementation, and adaptability under any subsequent evolutionary changes. In particular, any inherent complexities must be considered architecturally and either reduced or otherwise addressed constructively wherever possible, without compromising trustworthiness.

Background on various constructive approaches for coping with complex requirements and complex architectures can be found in [10, 13, 14] as well as [1, 2, 4, 16, 18, 19, 20], for example.

5 Conclusions

Myopia is very dangerous with respect to trustworthiness. The commonalities among different applications far beyond those considered here are likely to transcend any would-be single-discipline solutions. Thus, massive culture shifts are needed to consider information systems and their applications holistically and proactively. We must be able to develop systems and evaluate them in their entirety—especially through compositions of evaluated subsystems, with predictable aggregate behavior—and to ensure usability and the soundness of operational configurations. Of course, this culture shift is especially important for applications with critical requirements for trustworthiness. The pleas for such approaches in many seminal papers (including those that have been revisited previously in the ACSAC Classic Papers track) may seem old-fashioned, but are nevertheless still timely.

Inadequate understanding of the depths of the problems is also dangerous, in part because it typically leads to simplistic and untrustworthy solutions. Examples of that risk include beliefs that cryptographic certificates, longer passwords, firewalls, and testing can ensure security. However, too much ago-

nizing over the pervasiveness of these problems is also dangerous—because the resulting sense of hopelessness also typically leads to would-be solutions that are short-sighted and ineffective. In both cases, a danger of eschewing useful research and innovative developments, and placing trust in untrustworthy systems and people often prevails.

In particular, life-critical systems and other systems with stringent requirements for trustworthiness should be held to commensurately higher standards than conventional software, overcoming today’s realities—in which architectures are poorly structured, development practices generally yield numerous vulnerabilities, and criteria and proprietary evaluations are inherently incomplete.

Market forces appear inadequate in driving high-assurance components for critical systems. To be effective, open systems, open interfaces, and various forms of nonproprietary source code need more supporting incentives. Regulation is a slippery slope, as is attempting to invoke liability for the development or use of untrustworthy systems (which would entail some significant and difficult changes, including how to place the blame—as noted above). Insurance and tax incentives are possible, but also likely to have exploitable loopholes. Better awareness of the risks of untrustworthiness is clearly warranted. Above all, there are no simple solutions.

I am frequently asked why the ACM Risks Forum does not include more success stories. There are several reasons. Failures significantly outnumber successes. Success stories are rarely submitted, and I continue to hunt for them. Supposed successes are sometimes over-hyped, even masking development problems. I have cited Henry Petroski on numerous occasions; he noted long ago that we tend to learn little from successes, but that we have a better chance to learn from our failures. The RISKS experience suggests that we do not learn enough from either. Fortunately, a better understanding of the past seems to be emerging in some areas in recent years, along with perhaps fewer aircraft crashes, automobile recalls, and nuclear power disasters. The ACSAC classic papers track also provides some supporting perspectives—including those this year by Li Gong and Matt Bishop.

However, too many problems remain in areas such as health care, power distribution, malware, and elections, and lurking trustworthiness problems in automated highways, autonomous systems, and cloud computing. In any event, the lessons of past vulnera-

bilities, system development failures, and human limitations must be considered more pervasively, and a deeper understanding of the emerging threats and potential risks must be gained. I am optimistic that the research and development communities have much to offer, but am less optimistic that the needed culture shifts can be achieved because of the necessary major changes relating to governments, industry, education, economic policies, standards, procurement processes, entrenched interests, and so on.

Acknowledgments

This paper received support from ACCURATE: A Center for Correct, Usable, Reliable, Auditable, and Transparent Elections, under SRI’s National Science Foundation Grant Number 0524111. The author wishes to thank Douglas Maughan, who previously sponsored a program on high-assurance trustworthy systems (e.g., [10]) when he was a Program Manager in the Defense Advanced Research Projects Agency (DARPA), and who continues to be responsible for cybersecurity in the Science and Technology Directorate of the Department of Homeland Security, pursuing approaches aimed at preventing risks such as those mentioned here (e.g., [6]).

References

- [1] C. Boettcher, R. DeLong, J. Rushby, and W. Sifre. The MILS component integration approach to secure information sharing. In *27th AIAA/IEEE Digital Avionics Systems Conference*, St. Paul MN, October 2008. IEEE.
- [2] D.D. Clark and D.R. Wilson. A comparison of commercial and military computer security policies. In *Proceedings of the 1987 Symposium on Security and Privacy*, pages 184–194, Oakland, California, April 1987. IEEE Computer Society.
- [3] D.D. Clark et al. *Computers at Risk: Safe Computing in the Information Age*. National Research Council, National Academies Press, 2101 Constitution Ave., Washington, D.C., 5 December 1990. Final report of the System Security Study Committee.
- [4] E.W. Dijkstra. The structure of the THE multiprogramming system. *Communications of the ACM*, 11(5), May 1968.

- [5] S.E. Goodman and H.S. Lin, editors. *Toward a Safer and More Secure Cyberspace*. National Research Council, National Academies Press, 2101 Constitution Ave., Washington, D.C., 2007. Final report of the National Research Council Committee on Improving Cybersecurity Research in the United States.
- [6] D. Maughan et al. A roadmap for cybersecurity research. Technical report, Department of Homeland Security, October 2009.
- [7] R. Mercuri. *Electronic Vote Tabulation Checks and Balances*. PhD thesis, Department of Computer Science, University of Pennsylvania, 2001. <http://www.notablessoftware.com/evote.html>
- [8] P.G. Neumann. The role of motherhood in the pop art of system programming. In *Proceedings of the ACM Second Symposium on Operating Systems Principles, Princeton, New Jersey*, pages 13–18. ACM, October 1969. <http://www.multicians.org/pgn-motherhood.html>
- [9] P.G. Neumann. *Computer-Related Risks*. ACM Press, New York, and Addison-Wesley, Reading, Massachusetts, 1995.
- [10] P.G. Neumann. Principled assuredly trustworthy composable architectures. Technical report, Computer Science Laboratory, SRI International, Menlo Park, California, December 2004. <http://www.csl.sri.com/neumann/chats4.html>, .pdf, and .ps.
- [11] P.G. Neumann. Holistic systems. *ACM Software Engineering Notes*, 31(6):4–5, November 2006.
- [12] P.G. Neumann. Risks of untrustworthiness. In *Proceedings of the 22nd Annual Computer Security Applications Conference (ACSAC 2006), Classic Papers section*, Miami, Florida, December 2006. IEEE Computer Society.
- [13] P.G. Neumann. Reflections on system trustworthiness. In Marvin Zelkowitz, editor, *Advances in Computers, volume 70*, pages 269–310. Elsevier Inc., 2007.
- [14] P.G. Neumann. Hierarchies, lowerarchies, anarchies, and plutarchies: Historical perspectives of composable high-assurance architectures. In *Third Layered Assurance Workshop*, San Antonio CA, August 2009. AFRL. Slides: <http://www.csl.sri.com/neumann/law09+x4.pdf>
- [15] P.G. Neumann. Illustrative risks to the public in the use of computer systems and related technology, index to RISKS cases. Technical report, Computer Science Laboratory, SRI International, Menlo Park, California, 2009. Updated now and then: <http://www.csl.sri.com/neumann/illustrative.html>; also in .ps and .pdf form for printing in a denser format.
- [16] D.L. Parnas. On the criteria to be used in decomposing systems into modules. *Communications of the ACM*, 15(12), December 1972.
- [17] P. Porras. Reflections on conficker. *Communications of the ACM*, 52(10), October 2009. Inside Risks column. <http://www.csl.sri.com/neumann/insiderisks.html#219>
- [18] J.M. Rushby. The design and verification of secure systems. In *Proceedings of the Eighth ACM Symposium on Operating System Principles*, pages 12–21, Asilomar, California, December 1981. (ACM Operating Systems Review, 15(5)).
- [19] J.H. Saltzer and F. Kaashoek. *Principles of Computer System Design*. Morgan Kaufman, 2009. Chapters 1-6 only. Chapters 7-11 are online. <http://ocw.mit.edu/Saltzer-Kaashoek>
- [20] J.H. Saltzer and M.D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, September 1975.
- [21] M. Schaefer et al. *Multilevel Data Management Security*. Air Force Studies Board, National Research Council, National Academies Press, 1983. Final report of the 1982 Multilevel Data Management Security Committee.
- [22] F.B. Schneider and M. Blumenthal, editor. *Trust in Cyberspace*. National Research Council, National Academies Press, 2101 Constitution Ave., Washington, D.C., 1998. Final report of the National Research Council Committee on Information Trustworthiness.