Inside Risks 101 (*CACM 41,* 11, Nov 1998)
Towards Trustworthy Networked Information Systems
Fred B. Schneider

When today's networked information systems (NISs) perform badly or don't work at all, lives, liberty, and property can be put at risk.[1] Interrupting service can threaten lives and property; destroying information or changing it improperly can disrupt the work of governments and corporations; and disclosing secrets can embarrass people or hurt organizations.

For us—as individuals or a nation—to become dependent on NISs, we will want them to be *trustworthy.* That is, we will want them to be designed and implemented so that not only do they work but we have a basis to believe that they will work, despite environmental disruption, human user and operator errors, or attacks by hostile parties. Design and implementation errors must be avoided, eliminated, or the system must somehow compensate for them.

Today's NISs are not very trustworthy. A recent National Research Council CSTB study[2] investigated why and what can be done about it, observing:

• Little is known about the primary causes of NIS outages today or about how that might change in the future. Moreover, few people are likely to understand an entire NIS much less have an opportunity to study several, and consequently there is a remarkably poor understanding of what engineering practices actually contribute to NIS trustworthiness.

• Available knowledge and technologies for improving trustworthiness are limited and not widely deployed. Creating a broader range of choices and more robust tools for building trustworthy NISs is essential.

The study offers a detailed research agenda with hopes of advancing the current discussions about critical infrastructure protection from matters of policy, procedure, and consequences of vulnerabilities towards questions about the science and technology needed for implementing more-trustworthy NISs.

Why is it so difficult to build a trustworthy NIS? Beyond well known (to *RISKS* readers) difficulties associated with building any large computing system, are problems specifically associated with satisfying trustworthiness requirements. First, the transformation of informal characterizations of system-level trustworthiness requirements into precise requirements that can be imposed on system components is beyond the current state of the art. Second, employing "separation of concerns" and using only trustworthy components is not sufficient for building a trustworthy NIS—interconnections and interactions of components play a significant role in NIS trustworthiness.

One might be tempted to employ "separation of concerns" and hope to treat each of the aspects of trustworthiness (e.g., security, reliability, ease of use) in isolation. But the aspects interact, and care must be taken to ensure that one is not satisfied at the expense of another. Replication of components, for example, can enhance reliability but may complicate the operation of the system (ease of use) and may increase exposure to attack (security) due to the larger number of sites and the vulnerabilities implicit in the protocols to coordinate them. Thus, research aimed at enhancing specific individual aspects of trustworthiness courts irrelevance. And research that is bound by existing subfield demarcations can actually contribute more to the trustworthiness problem than to its solution.

Economics dictates the use of commercial off the shelf (COTS) components wherever possible in building an NIS, which means that system developers have neither control nor detailed information about many of their system's components. Economics also increasingly dictates the use of system components whose functionality can be changed remotely and while the system is running. These trends create needs for new science and technology. For example, the substantial COTS makeup of an NIS, the use of extensible components, the expectation of growth by accretion, and the likely absence of centralized control, trust, or authority, demand a new look at security: risk mitigation rather than risk avoidance, add-on technologies and defense in depth, and relocation of vulnerabilities rather than their elimination.

Today's systems could surely be improved by using what is already known. But, according to the CSTB study, doing only that will not be enough. We lack the necessary science and technology base for building NISs that are sufficiently trustworthy for controlling critical infrastructures. Therefore, the message of the CSTB study is a research agenda and technical justifications for studying those topics.

Cornell University Professor **Fred B. Schneider** chaired the CSTB study discussed in this column.

[1] See Inside Risks ("Protecting the Infrastructures") in *CACM 41, 1* for a summary of the President's Commission on Critical Infrastructure Protection (PCCIP) report which discusses the dependence of communication, finance, energy distribution, and transportation on NISs.

[2] The final report for the Information Systems Trustworthiness study by the Computer Science and Telecommunications Board (CSTB) of the National Research Council, *Trust in Cyberspace,* can be accessed at http://www2.nas.edu/cstbweb/index.html.