High Confidence Systems Working Group

Committee on Computing, Information, and Communications Research and Development

National Science and Technology Council

SETTING AN INTERAGENCY HIGH CONFIDENCE SYSTEMS (HCS) RESEARCH AGENDA:

Proceedings of the Interagency High Confidence Systems Workshop

25 March 1998

Hosted by the National Coordination Office for Computing, Information, and Communications
Arlington, Virginia

I. INTRODUCTION

On 25 March 1998, the High Confidence Systems (HCS) Working Group, under the auspices of the National Science and Technology Council's Committee on Computing, Information, and Communications (CCIC) Research and Development (R&D), sponsored an invitational workshop to begin setting an interagency HCS research agenda that can be used to help establish new Federal research funding initiatives. The workshop was hosted by the National Coordination Office for Computing, Information, and Communications (NCO for CIC), Arlington, Virginia, and chaired by Ms. Teresa Lunt, Program Manager, Information Survivability, Information Technology Office/Defense Advanced Research Projects Agency.

The objectives of the workshop were to:

- Elicit agency perspectives on HCS
- Establish a set of challenging research goals that could help drive the research and support the necessary funding
- Continue to build the critical mass of HCS research support

To achieve its objectives, the workshop brought together individuals from various agencies possessing HCS issues, garnered their perspectives, and began the process of setting the HCS research agenda by setting forth a set of top level goals.

BACKGROUND

The 1995 CCIC workshop and 1997 HCS workshop[1] explored HCS issues, needs, potential solutions, and future technology research that could improve the state of safety- and security-critical systems. From the two workshops, the clear need for developing a national agenda for HCS research emerged. Further, the 1997 workshop made it clear that the timing was right for pursuing such research, in part because of technology advances and also because the projected performance goals of various agencies were becoming dependent upon achieving improvements in HCS technology.

WORKSHOP ORGANIZATION

The 1998 HCS workshop was conducted in three parts: Review of the 1997 HCS Workshop, Agency HCS-Related Research Program Overviews, and HCS Research Agenda Development. The first session brought all of the attendees up to date on the results from the previous workshop. The second session allowed each agency to express its HCS research needs, desires, and plans. The final session consisted of a group brainstorming effort to coalesce agencies' needs and desires into a set of goals that would support an interagency HCS research agenda.

II. REVIEW OF THE PRIOR HCS WORKSHOP

The Workshop Chair began this workshop by reviewing what took place at the last HCS Workshop on 6-7 August 1997—a draft of the proceedings was handed out. In her review, she summarized the key workshop ideas and noted some significant trends:

---

[1] *America in the Age of Information: A Forum*, July 1995, is available at http://www.ccic.gov. *Research Challenges in High Confidence Systems*, August 1997, will also become available at that URL.

- There is a growing dependence on computing for safety-critical use.
- Many of these safety-critical products are consumer hardware and/or software.
- There is an increased risk due to integrating functions that were formerly separated for safety purposes.
- There is increased computer use for improving system functionality.
- Infrastructures are running at capacity or near capacity.
- There are many high-visibility problems with complex systems.

Today, more than ever, the U.S. needs high confidence systems. A national HCS research agenda is needed *now!* Such an agenda must focus on safety and security in malicious environments because current techniques may be reaching their limits in systems of increasing scale and complexity. Recent technology advances put us on the brink of new possibilities; however, the U.S. is being overtaken by Europe in this area. To keep the U.S. in the forefront and to maintain our economic competitiveness, the U.S. must pursue advancing this technology for its own systems. Creating a national HCS research program will require a future-looking view of the world with outcome-oriented goals that appeal to the public, are affordable, do not generate fear of technology, and have a technology strategy.

Recommended HCS solutions should encourage lower cost, the leveraging of knowledge within a company, and techniques that do not require large educational ramp-up. While it was noted that many companies have a strong corporate incentive not to share their proprietary tools, the hope is that increasing integration will create a need for companies to share information on assurance techniques.

Part of the discussion that occurred during the review presentation focused on the role of standards, although that role is perceived differently by various participants. Some participants felt that standards are needed to focus research funding and to drive technology research (e.g., performance or safety standards that future HCS technology must meet). Others felt that research should be driving, or even developing, technology standards (e.g., standardized application programming interfaces (APIs), standardized security mechanisms) or interoperability standards (e.g., standardized certification authorities). Still others believed that standards only identify what is known and that research should look beyond standards to the future; the focus should be on the research agenda and letting standards developers adjust accordingly rather than on standards that force research adjustments. This third group felt that research could be putting HCS fundamentals in place to support current and future standards. While all positions appeared to have merit, no consensus was reached regarding how best to stress standards in this research agenda effort. Emerging from this discussion was an increasing need to move from process control-based standards more towards product-based standards for high confidence systems. There was also discussion about the Federal Highway Administration's Intelligent Transportation Systems (ITS) Program, which is now coming out with critically needed standards (including standards for software).

III. AGENCY HCS-RELATED RESEARCH PROGRAM OVERVIEWS

(1) Federal Aviation Administration

An overview of the Federal Aviation Administration's (FAA) far-reaching (10-20 years) research plan was presented to stimulate the workshop participants' thinking about future

HCS research. The National Airspace System (NAS) is highly complex and probably the largest system of systems in the world. The overview began with a description of the role of high confidence in air traffic control and the activities taking place in a NAS working group. Results from HCS research are expected to help meet many user needs including more airspace flight capacity, improved procedural efficiency, shorter travel times, increased safety. It was noted that introducing new technical capabilities can introduce vulnerabilities (e.g., GPS susceptibility to jamming). Such new technical capabilities, when provided with open systems using commercial-off-the-shelf (COTS) software, almost certainly will introduce vulnerabilities. These vulnerabilities, if exploited, will likely possess operational impacts such as lost time, wasted fuel, or decreased safety margins. Thus, there is a need to introduce security services into the NAS as well as incorporating the traditional architectural alternatives such as redundant systems (e.g., Wide Area Augmentation Systems backup technology to GPS).

The FAA NAS working group has on-going work in architectural approaches that focus on both safety-critical portions of systems and reducing costs. They are working on testing and verification of software integrity, and streamlining software aspects of certification to achieve faster certification at reduced costs with higher confidence in system safety. The current certification takes at least 3 years and is process intensive. The FAA would like the process to become more product oriented and needs future high confidence research to address shortened product evaluation times (e.g., three to four months).

In the information security (INFOSEC) area, the FAA needs future research to develop and mature a strategy for securing a system of multiply interconnected systems. Their work is categorized under four broad areas: (1) policy application, (2) INFOSEC implementation, (3) intrusion detection, and (4) standards. Two of these areas were addressed in more detail.

- In the area of policy application, they are working on security policy development that, in turn, is leading to a concept of operations driven by security. This work supports an evolving logical security architecture (currently version 3.0) to support distributed applications. They need further research that will enable them to implement a wider variety of security policies and to translate system administration policies to resources.
- In the area of INFOSEC implementation, they are working on assessment of vulnerabilities and risks, and development of security architectures and standards. HCS research must promote secure, interoperable, heterogeneous computing systems. Such research should also support a life-cycle protection strategy. Specific research must aid in layering access control mechanisms over operating systems and provide capabilities to authenticate network transactions for integrity. FAA NAS also needs more work on assessing its system vulnerabilities and associated risks. FAA NAS is currently focused on outsiders, but will include insiders.[2] FAA NAS must be able to deal with a disgruntled insider while protecting employee privacy. A desirable approach for addressing this issue is to segregate the NAS into communities of interest to limit the damage that any insider can do (soft-failure strategy).

(2) Federal Railroad Administration

The FRA's viewpoint of HCS focuses on the need for positive train control (PTC) to enforce train movement and speed limits and to reduce the probability of collisions. Existing railroad signal systems are extremely reliable, but most still permit one person to make a mistake that causes an accident (e.g., less than 5 percent of today's railroads have

[2] Statistics show that 70 percent of threat comes from the inside.

any automated enforcement of signal indications). PTC leaves people in the loop, but intervenes in an automated fashion if the people (e.g., engineers, dispatchers) do not respond properly. PTC would use differential GPS and dead reckoning to determine the location of trains and maintenance-of-way equipment.

Published and unpublished industry studies have shown that with PTC, the probability of collisions and overspeed accidents would be lowered by a factor of 100 and the annual rate of return on the investment would be 30 percent. In 1993, however, railroad company Chief Executive Officers terminated the industry's PTC program. Possible reasons for this termination may include: concerns about government regulation, decisions to invest in mergers rather than technology, concerns about the estimated capital investment costs ($3-$4 billion for all U.S. railroads), fears of liability from acknowledging that PTC is safer than current train control systems, and uncertainties about the effectiveness of the technology. In 1994, the FRA submitted a report on PTC to Congress that indicated FRA would initiate a regulatory process for PTC in FY 1997. The collaborative rule making has begun, but is moving very slowly.

FRA has a $20 million research budget, all devoted to safety-related projects. When FRA requested additional money last year to pursue PTC, the Office of Management and Budget (OMB) denied their request. This HCS effort may assist in getting some funding to further pursue PTC.

(3) Federal Transit Administration

The FTA is one of the nine operating administrations or agencies of the U.S. Department of Transportation and carries out the Federal mandate to improve public mass transportation. FTA is the principal source of Federal financial assistance for the planning, development, and improvement of public transportation systems. The Office of Research, Demonstration and Innovation of the FTA, in consultation with other government agencies and the transit industry, initiates projects aimed at improving mobility, economic growth and trade, safety and security, and human and natural environment. Activities include research, testing, evaluation and documentation, deployment, standards/architecture development, and mainstreaming/implementation. Because the research budget is small, FTA acts as a catalyst and driving force to leverage small research investments for enhanced results in identified priority areas.

At present, rail transit agencies are looking for a reliable and cost-effective solution to improve train throughput, using existing infrastructure, and to concurrently obtain improved safety of several orders of magnitude. Communication-Based Train Control systems that use modern communications, control, and computer technologies, offer a viable solution. As with other processor-based systems for safety-critical, real-time applications, there is a pressing need to develop methodology for safety verification and validation that is relevant to the application. Techniques such as numerical assurance, checked-redundancy, N-version programming, and diversity and self-checking are being considered. FTA is also supporting development of standards through the Institute of Electrical and Electronic Engineers, a standards developing organization to promote commonality in functionality, operations, and interoperability.

(4) Food and Drug Administration

The FDA is not a research funding organization. With respect to HCS, there is a $2 million budget that supports one of the five FDA centers looking at product approval issues. This center is concerned with software in two areas: production of products and process control. FDA deals with a wide range of manufacturers (median size is 50 employees, and none of

the companies shares information) and a wide range of user expertise (doctors in a hospital to patients at home), which leads to user interface problems. Since the FDA does not have much money for research, they focus on supportive efforts (e.g., education, guidance, and standards).

Given the pressure to use COTS software, the FDA is interested in research on how to validate COTS software. There is concern about near-term solutions because verification is currently left to the vendor of each product. It would be helpful to have evaluation methods to run existing products through a standard test suite. Legally, the FDA can only regulate products marketed commercially—they cannot regulate what doctors privately develop and use as part of their practice of medicine within their own offices.

(5) National Library of Medicine

The National Institutes of Health (NIH), of which NLM is a part, is most interested in research on issues related to High End Computing and Communications (HECC) and the Next Generation Internet (NGI). Medicine and health care are more probabilistic than high confidence. NIH is not developing HCS technology, but uses high confidence systems and pushes them to the limit. The greatest need within NIH and the practice of medicine is for confidentiality, integrity, and availability of health records. The greatest security threat is insiders. Their goal is to achieve 100 percent availability with appropriate access control—they cannot "fail safe" because a lack of information could mean that a patient dies. They also need reliable computation systems (e.g., for image and sound enhancement) and control systems (e.g., for telemedicine and telesurgery).

Current technology research funding goes toward demonstrating technologies that are given to NIH. A small amount of HCS-related research is addressing wide area networking for hospitals to integrate their functions.

(6) Department of Treasury

The Department of Treasury is a very diverse organization with diverse missions supported by systems that must survive in all kinds of environments and be trusted. They have $1.8 billion for information systems research—$1.2 billion goes to the Social Security Administration, and there are no direct funds for HCS. Treasury conducts red teams and tries to require vendors to use the Capability Maturity Model for software development. They have to accommodate a variety of technologies and user capabilities when they deliver services to customers (e.g., many do not have access to a computer, bank account, or even a telephone).

Treasury's problem areas include engineering process integration, requirements definition, user/customer interface, and programming tools and techniques. Treasury has big legacy programs that continue to work on a quagmire of technologies, but they are trying to look out beyond ten years to a whole interface and interoperability. Their most important requirement is security/safety of the services they provide. They have some pilot projects with banks (e.g., smart cards), and they are always looking to improve security for physical and electronic transactions while keeping an audit trail of accountability. They are wary of partnerships in industry where products become co-dependent, tying organizations to particular products.

(7) Secret Service

The two primary goals of the Secret Service are to be the premier law enforcement agency and to protect financial transactions. The Secret Service gets involved in a wide variety of

research (e.g., behavioral, engraving/printing, chem/bio). They have a pressing need to facilitate communication between different groups (e.g., Secret Service, Park Police, city police, ambulance). Within the National Performance Review (NPR) Information Technology Objectives, ITO-4 contains a research project looking at a wireless infrastructure for all law enforcement and public safety communities. One particular concern within this wireless communications infrastructure is being able to reserve part of the communications spectrum for law enforcement.

(8) National Aeronautics and Space Administration

NASA provided a program overview at the last HCS Workshop. Updating that overview, there was brief mention of an evolving NASA program to develop a safe, robust, secure datalink to assure free-flight avionics systems. It was noted that there will be a workshop at the end of April at NASA-Ames to define the program. The update focused primarily on aeronautics (rather than on space). There is an Aviation Safety Investment Strategy Team (ASIST) that consists of over 200 industry representatives plus some government people. Four of the primary investment areas that were identified by the ASIST subteam on Flight Critical Systems and Information Integrity (FCSII) were discussed:

- *Software certification:* NASA-Langley is leading a short-term effort called Streamlining Software Aspects of Certification to reduce validation and verification costs.
- *Analytical investigations of flight critical/essential systems:* Certifications based on analysis of product rather than control of life cycle (i.e., a methodology based on formal methods because reliability assessment is infeasible). Process control will not be thrown away, but will be streamlined.
- *Pilot-vehicle interactions:* Multi-disciplinary research involving both human factors and formal methods.
- *Robust partitioning:* The aviation industry wants an avionics computer resource (ACR) that could execute multiple applications with different levels of criticality on a single processing site that has been pre-certified by the FAA. Mechanisms for assuring logical partitioning in a real-time, fault-tolerant system are not yet developed, and a standard application program interface is needed (the marketplace is unlikely to produce such an API). NASA-Langley is developing fundamental formal models of partitioning and is formally verifying Honeywell's SAFEbus.

NASA will spend $500 million over five years for aviation safety — $12 million per year will go toward information systems.

(9) Department of Energy

DoE also provided a program overview at the last HCS Workshop. Updating that overview, the discussion focused on DoE's "open" side (as opposed to the "dark" or defense side), where they do a lot of work with international groups and universities. Two projects that are currently on going at DoE were described:

- *DoE 2000:* Started two years ago, this project looks at collaborative environments, protection of intellectual property, and *Akenti* (usage-based authenticated access control system requiring certificates with a list of capabilities).
- *Accelerated Stockpile Computing Initiative (ASCI):* Nuclear weapons stockpile protection, using distributed simulation and needing ultra-high speed cryptography.

DoE is funding work in incident response, information warfare, intrusion detection, high speed cryptography, distributed computing, secure software distribution, and reliable/ordered multicast. The major issues at DoE are key escrow (industry), export control, very high-speed cryptography, multilevel/policy authentication and access control, public-key infrastructure (PKI) resiliency, cross PKI/domain operation, and striping of communications (efficiency and fault tolerance). It was noted that all agencies need to be involved in researching the infrastructure for passing public-key certificates.

(10) Nuclear Regulatory Commission

The NRC update of its previous overview began with an explanation of NRC's standard plan for reviewing software, which focuses mostly on process but is also looking at the product. There is a new regulation (NUREG/CR 6463), available at www.nrc.gov, that identifies computer language features *not* advisable to use in high confidence systems.

The NRC believes they will achieve a high payoff by looking at hazards from a system view. The next phase of this research will be to decompose system requirements into system, component, and human parts. They are also doing research in software reliability metrics—they have found that different metrics have different benefits at different times during development. There will be a metrics conference in Bethesda, MD, in November 1998.

Other NRC HCS research efforts include working with the University of Virginia on modeling and simulation, addressing sampling rates of digital systems, and looking at various tools to evaluate products. Risk-based, performance-based evaluation is a new research area being pursued by the NRC because no one knows how to prove high reliability (e.g., 10^{-9}).

(11) National Institute of Standards and Technology

The NIST update began by noting that until recently, NIST activities have been driven by the Brooks Act. Now, NIST is trying to work with industry to make products better through measurement and test of software. Approximately $2.5 million is devoted to this effort, which has three thrusts:

- *Development methods and tools:* Developed a program slicing tool for C (developing one for Java) and examining role-based access control (good for medical informatics and electronic commerce).
- *Data collection and analysis:* Failures in different kinds of software and methods used to detect those failures—promoting sharing among vendors with similar concerns and developing a standard for comparing tools and goals.
- *Specification and testing:* Most interested in formal methods, but industry is reluctant —NIST feels they can make formal methods more interesting to industry by developing formal test methods/tools to reduce product testing costs (automated testing based on formal specification).

(12) National Institutes of Health

The NIH representative noted in his update that a Presidential Commission on quality in the health care industry has also documented the shock statements he previously provided at the August 1997 workshop. Having compared standards with research budgets—standards

groups are using standards to drive research being performed by companies in that area¾ the NIH representative recommended research be focused on the critical portions of *real-world* systems. These sorts of workshops help to diffuse HCS research, which is vital to the NIH community. It was noted in the update that the medical informatics community wants no regulation on their work, but wants heavy regulation on anything from outside their community. Also mentioned, was a 30-year old study on uncertainties and the need to focus on them.

IV. HCS RESEARCH AGENDA DEVELOPMENT

This brainstorming session began with additional context setting. The NCO representative provided organizational background information on the CIC HCS effort and where it fits in with other CIC Program Component Areas. It was suggested that, in the course of setting the HCS research agenda, a document like the NGI Implementation Plan (NGI IP) be produced (the NGI IP took a little over a year to complete). Proposed research issues that might be included in such a document were mentioned (e.g., open systems, cost reduction, integration across technologies and applications) along with organizational/operational issues related to the HCS research agenda (e.g., role of the Federal government, individual agency roles, multi-agency coordination).

To further set the context, a short presentation on NSA's perspective of researching high confidence in INFOSEC systems was provided. NSA has moved away from the *old world* Government-off-the-shelf (GOTS) approach to assurance (i.e., government control, lots of effort placed on assurance because systems are in the field for a long time, high quality secure products). It is moving towards the *new world* commercial-off-the-shelf (COTS) approach to assurance (i.e., government has no control, little effort is invested in assurance because of short product life cycles, layered security). The main problems with assurance are finding and keeping experts and needing tools that have assurance built into them. Five "Pillars for Success" were identified as potential starting points for our brainstorming (fundamental property research, design tools, administration tools, composition, and education). Additionally, six "Stretch Goals" (reduce credit card interest, safeguard children in cyberspace, privacy of consumer transactions, decrease system intrusions, eliminate junk email, and dominate the electronic battlespace) were identified.

The Executive Secretary of the HCS Working Group proposed a strawman HCS vision and a set of technology research goals for 2010:

VISION

High Confidence Systems research will produce a body of knowledge and a set of tools that will promote safe, reliable, dependable, secure, and survivable computing systems. Achieving high confidence systems will enable a world radically different and much improved from the world of today in providing for the general welfare of the public and meeting key performance goals of Federal agencies.

HCS Technology Research 2010 Goals

- Develop the foundational capabilities to specify and ensure implementation of requisite behaviors in highly complex, large-scale systems
- Develop techniques to measure system qualities
- Develop techniques to incorporate user-centered needs

- Develop and integrate tools and techniques that support HCS implementation and assessment
- Validate that HCS tools and techniques scale to real problems
- Lower the costs of HCS implementation and assessment
- Integrate HCS approaches more widely into all government and commercial system development environments

To continue setting the context, a NASA folder (illustrating the results of NASA's efforts to shape their research agenda) was distributed as an example for getting an HCS research agenda established. A strawman strategic roadmap for HCS with five "pillars" (reliability and fault tolerance, security, survivability, safety, and assurance technologies) was proposed to get the group started (see Appendix B).

The brainstorming began with the moderator pointing out that the group needed to get to clear, concise goals established to convey the research agenda to the Administration, Congress, and the public in a way that all could understand almost immediately. With that perspective, it was suggested that the group identify the problem(s) first, in a way that the average person can understand (e.g., delivery of government services); otherwise, the group's output would be a solution chasing a problem. Once again, there was some discussion about the role of research related to standards. There also was a question about the scope of the research agenda:

- Government's role in national security
- Government's role to itself and to private industry
- Private industry's role unto itself

It was suggested that other groups are already handling the first and last items of scope. The scope of the interagency HCS research agenda should focus on Government's role to itself (efficiency of Government services) and to private industry (items of public interest, leading standards, and working issues that industry will not address on its own). There was an open question about being able to influence OMB, which is making a lot of decisions for the various agencies, in its apparent reluctance to support needed research.

A number or slides were developed as the group brainstormed ideas and how to structure them. Three slides were developed at the end of the afternoon to capture the group's thoughts at that point on high level goals that might convey the importance of HCS:

1. PROTECT THE PUBLIC: Increase confidence in critical infrastructures.
   - "Confidence" includes safe, reliable, accurate, trustworthy, secure, adaptable, and timely.
   - "Critical infrastructures" includes medical, transportation (e.g., aviation, trains), power, telecommunications (e.g., GPS radio navigation), public safety and emergency services, national security and defense, command and control, financial services, and the environment.
   - Measurement, design analysis, and formal methods are important.

2. PROTECT THE CONSUMER: Enable higher reliability and ease-of-use in commercial products.

- "Reliability and ease-of-use" includes expedited certification, validation & verification, shortened time to market, simplicity of use, affordability, lower life cycle cost, and plug & play.
- "Commercial products" includes smart cars (Intelligent Vehicle Systems), medical devices, consumer electronics, business systems, smart houses, sensor technologies (e.g., alarms), GPS receivers, consumer Internet, smart cards, education technologies, digital libraries.
- Informal test and evaluation are important.

3. ENHANCE GOVERNMENT SERVICES: Increase confidence in Government services.
   - "Confidence" includes timely, accurate, private, adaptable, safe, secure, dependable, accountable, responsive, easy to use, and efficient.
   - "Government services" includes entitlement programs, IRS modernization, Social Security modernization, Medicare modernization, law enforcement modernization, emergency management modernization, Patent and Trademark Office (PTO), air traffic control, and NOAA.

These three slides were delivered to all attendees who were tasked to provide feedback. A smaller group will examine the feedback and create a polished set of ideas for discussion at a future workshop. The goal is to continue the development of a document for HCS research similar to the NGI IP, and to publish a first draft of that document by the end of this summer.

The Workshop Chair closed the meeting by thanking everyone for their contributions and by encouraging increased participation in the HCS Working Group.

APPENDIX A. WORKSHOP PARTICIPANTS

Mike Ackerman, NLM
Bob Aiken, DoE
Robert Brill, NRC
Wayne Bryant, NASA
Ricky Butler, NASA
Tice DeYoung, NASA
Steve Ditmeyer, FRA
John Faust, AFRL
Helen Gill, DARPA
Henry Heffernan, NIH
Sally Howe, NCO
Kay Howell, NCO
Ron Kangas, FTA
Feisal Keblawi, FAA
John Kinkel, FAA
Rick Kuhn, NIST
Jerry Linn, NIST
Teresa Lunt, DARPA
Terry Mayfield, IDA
Robert Meushaw, NSA
John A. Murray, Dept. of Treasury
John F. Murray, FDA
Mike Nassif, AFRL
Rob Rosenthal, NIST
Harvey Rudolph, FDA
Brian Snow, NSA
Ron Thomsen, Secret Service
Dolores Wallace, NIST
Steve Welke, IDA

APPENDIX B. NASA STRATEGIC ROADMAP STRAWMAN