# History of US Government Investments in Cybersecurity Research

A Personal Perspective

Carl E. Landwehr
Institute for Systems Research
University of Maryland
College Park, USA
Landwehr@isr.umd.edu

***Abstract*—This paper traces the history of cybersecurity research funding by the U.S. government. Difficulties in accurately measuring the level of U.S. government research funding for cyber security are first described. Some of the legislative and bureaucratic mechanisms involved in funding and reporting such research today are reviewed. A qualitative, personal perspective on the ups and downs of US cybersecurity research funding from the late 1960s to 2010 is then provided. The essay is written for the thirtieth anniversary meeting of the IEEE Symposium on Security and Privacy, held in May 2010.**

***Keywords-cybersecurity research; research funding; HPC; NCO; NITRD; CSIA; information assurance; computer security; information security***

## I. Difficulties in Measuring Government Funding for Cybersecurity Research

Government funding over the past 30 plus years for what is now called cybersecurity research (also known over this period as computer security, information security, and information assurance research) has played a fundamental role in advancing knowledge of how to build components and systems that can meet their specifications and behave as users intend, even in the face of malicious behavior. Both the substance of that knowledge and the reasons much of it has lain unused in actual system developments are topics for another essay, but the fact that many ideas developed in the early years of cybersecurity research seem now to be resurfacing provides some reassurance that the work funded by the early investments has had lasting value.

The annual IEEE Symposium on Security and Privacy is the product of continuing government research funding in this domain. Perhaps someone can quantify the fraction of the research reported in the Symposium that was not directly stimulated by government research investment; I suspect it would certainly be under half, very likely under a quarter, and possibly under a tenth. The level of government research funding has fluctuated over the years and with it, sometimes, the level of participation in the symposium.

So, can we characterize over the past 30 years, in rough terms, the level of US government investment in cybersecurity research? Some questions of interest to taxpayers (and hence Congress) might be:

- When did the U.S. government start research investment in this area?
- How much has been invested, by year and by agency?
- What results have come from this investment?

Answering the third question could be a lengthy and possibly contentious exercise and is beyond the scope of this paper. The first two questions might seem to be relatively simple accounting questions, though, and indeed they are the kinds of questions for which the U.S. Congress likes to have clear answers. Here are some reasons why they are quite hard to answer.

First, cybersecurity is a broad and complex topic. "Secure" behaves as an adjective and it can modify many nouns. Hence there is research in secure hardware, operating systems, programming languages, networks, and user interfaces, not to mention in specifying security requirements, in tools of all sorts for developing systems that can meet those requirements, and in assessing the extent to which the developed system really does meet those requirements. Further, there is research in detecting when systems have been compromised, recovering from compromises, and forensics to determine how the compromise occurred and who should be held accountable for it. One of the fundamental technologies used to provide both security and privacy is cryptography, which probes deep into the heart of the theory of computer science. More recently, research in program analysis has been seen to have substantial applications in identifying security vulnerabilities. In short, a substantial part of computer science and engineering can be looked at through the lens of security and privacy. It is impossible to draw a sharp line between the funding in each category, and often a single research project may contribute to both categories.

Second, the way the U.S. government authorizes, appropriates, and expends money is a complex process in which many different parties each play different roles and exert influence on the way research dollars are spent. Typically, for each department and agency, Congress passes an authorization bill that lays out in some detail limits on what an agency can spend in a particular area. The authorization does not provide actual dollars, only the authority to spend them. The actual dollars come from appropriation bills that may or may not reflect the authorizations that have been passed. Most often, the authorizations will be higher than the appropriations, requiring the agency to choose which authorizations it will fund more fully. So one cannot simply look at what Congress authorizes and determine how much has been spent on what. The appropriations are more informative, but they still often leave significant discretion to the agency receiving the funds

Further, the government funds research in many different ways and through many spigots. The most visible expenditures to the general public are those awarded through public and open competitions such as the solicitations issued by NSF, DARPA, IARPA, DHS, ONR, AFOSR, ARO, etc. But some very fruitful research is also conducted within government laboratories that receive their funding through different appropriations and which have their own mechanisms for distributing those funds internally.

Finally, some cybersecurity research is conducted out of the public view. Budgets for classified research are usually not publicly available, at least in detail.

In the early 1990s, Congress decided to fund an initiative in High Performance Computing (HPC) but was concerned that many agencies were, in the way we've just seen, funding related efforts in an uncoordinated, or at least unreported, fashion. It set up a National Coordination Office (NCO) as an explicit mechanism to orchestrate interagency cooperation and coordination on this topic. Subsequently, the NCO has developed interagency working groups in several areas beyond HPC, covering what is now known as the Networking and Information Technology Research and Development (NITRD) program. The most recently created of these is the Cyber Security and Information Assurance (CSIA) Interagency Working Group (IWG), chartered in August 2005. It meets monthly and includes representatives of many (but not all) Federal departments and agencies that fund R&D in cybersecurity. It publishes an annual supplement to the President's budget that includes cybersecurity research funding requests by each of these agencies.

Table 1 summarizes the figures reported by the CSIA IWG since its inception [1]. Note that the Office of the Secretary of Defense (OSD) and Department of Defense (DOD) figures for the first year are artificially low, that only unclassified research is covered, and that not all research funding agencies are included (notably, Department of Homeland Security (DHS) and Intelligence Community funding is not reported). The Department of Energy (including the National Nuclear Safety Administration) and other NITRD agencies not shown did not report any spending on CSIA research until FY10. The estimated FY10 actual spending figures in Table 1 are draft, subject to final corrections.

## II. A PERSONAL PERSPECTIVE ON THE EBB AND FLOW OF RESEARCH FUNDING

With this explanation of why it's hard to assess precisely how much the government has funded in cybersecurity research as background, this section provides a subjective and longer term view of overall trends in the level and significance of cybersecurity research funding.

U.S. government funding for computer security research had begun by at least the last half of the 1960s with support from ARPA (as today’s DARPA was then known) for the ADEPT-50 system, which achieved operational use. Built with IBM System/360 hardware, it was based on a formally expressed security model. Also in the late 1960’s, Multics development was initiated under ARPA sponsorship for Project MAC.

In 1970, a Defense Science Board (DSB) task force chaired by Willis Ware of the RAND corporation issued an influential report (subsequently known as Ware report [2]). Experience gained in building (and penetrating) ADEPT-50, as well as Multics design experience contributed to that report, which was classified Confidential at the time because it could reveal government policy, not because it contained any sensitive technical information. The report addresses computer security in the context of resource-sharing (including time-sharing) computer systems, information of different classification levels, and users with different clearances. The memorandum conveying the report to the DSB states: "It is important to influence designers of future computers and software so that security controls can be installed before the fact and as an integral part of the system … Thus a program of studies and research is required. This need should be made known to various agencies of the Department of Defense … some aspects of the program are appropriate for ARPA."

Multics provided a source of empirical demonstrations for the graduate work of Roger Schell, who contributed more than 10,000 instructions to the core OS as a student member of the development staff. After graduating, (then) Major Schell, Ph.D., was assigned to ESD Hanscom and given responsibility for two early-stage security projects. One was to enhance Multics security for multilevel operation in order to replace less secure GE 635 GCOS systems then installed in the Pentagon, enabling Multics to process both Secret and Top Secret data. The other was to contract for and manage the panel of experts that produced the Anderson report [3], whose scope was to "develop a comprehensive plan for research and development leading to the satisfaction of requirements for multi-user open computer systems which process various levels of classified and unclassified information simultaneously through terminals in both secure and insecure areas." The report recognized a need for a formal definition of what is meant by a secure system and advocated systems built around a security kernel incorporating a reference monitor.

The Ware and Anderson reports stimulated research funding that issued primarily from ARPA, but also from the Air Force, Navy, and Army. Both reports are still worth reading and are freely available.

Major Schell managed to find funds to support a portfolio of security projects over several years despite the reluctance of the Air Force bureaucracy to admit there might be a need for this work. Each year, Air Force research administrators would strip funds for computer security research, which focused on varous aspects of providing mulitilevel security for networked, time-shared systems, from the budget. Apparently, some people involved came from a background in nuclear weapons modeling and believed that computers worked more efficiently when dedicated to one problem at a time. Since sharing and networking computers was inefficient and wasteful in the first place, any research encouraging such use (e.g., by permitting multilevel operations) would simply encourage inefficiency. Nevertheless, each year Major Schell would manage to find unspent money from other projects late in the fiscal year ("sweep up" funds), often from operational organizations who saw a specfic need for multilevel operation, rather than from research funders, to keep these projects afloat.

Much of this work initially focused on the development of security kernels, what policies they might enforce, and how they might be used in practice. A design for a security kernel for Multics was funded, for example (although never implemented -- security improvements for Multics were implemented but the kernel was not). In the late 1970's, the Air Force bureaucracy finally triumphed when a Brigadier General was persuaded to issue a "stop work" order on all of its computer security research contracts. The Army and Navy also funded research projects at this time (and provided some of the "sweep up" money that had found its way into the Air Force portfolio).

In the late 1970s, Steve Walker left his position at NSA to become a program manager at ARPA. There he managed several computer security related projects, including a kernelized VM/370 retrofit (KVM/370) and security kernel developments by Ford Aerospace (KSOS) and Honeywell (SCOMP). The SCOMP project had been initiated by Roger Schell as part of the Air Force ADP Security Research Program to provide a multilevel communications front end for Mutlics. KSOS was a separate effort to implement an industrial-strength security kernel for a minicomputer (PDP-11); each of the efforts eventually found it needed to provide a Unix-compatible system call interface in order to gain acceptance. ARPA did all of its contracting work through other agencies (this approach was intended to help keep the military informed about what ARPA was doing and to facilitate technology transfer). The KSOS work was funded through NSA, and the SCOMP work was funded first by the Air Force and later through the Navy. The services did also fund some computer security research independently, but it was not a major area of funding for any of them. It is worth noting that this early work also involved investment in tools for the formal specification of software and for proving properties of those specifications.

Steve Walker subsequently left ARPA for the Pentagon and began to try to get more resources devoted to the area. Under the rubric of the "Computer Security Initiative" he convened regular meetings of representatives from all the services and (if memory serves) NSA in the Pentagon as a means of coordinating the independent research funding efforts and building momentum for greater support. One difficulty in funding security research at the time was that the funding had to compete with research targeted at developing new weapon systems (tanks, ships, airplanes), and it was difficult to argue that computer security deserved the same degree of attention as research leading to a new tank.

Funding for many of the projects Steve initiated at ARPA continued after his move to the Pentagon. With Steve's help, KVM/370 gained support from the US Army and Canadian Department of National Defense; international financing arrangements were facilitated by the purchase of a virtual canon (or was it a tank?).

Steve was persistent and effective. He pursued the development of a method for grading the security provided by systems and of supporting system security evaluations using government resources. This effort led to the establishment of the DOD Computer Security Evaluation Center (quickly shortened to DOD Computer Security Center; a few years later it was rechristened the National Computer Security Center (NCSC)). The director of NSA, ADM Bobby Inman, easily won a bureaucratic duel with NIST over the location of the center. Initially, however, the NCSC was a unique institution at NSA in that it was relatively open about what it did and the products and research it produced. The Trusted Computer System Evaluation Criteria (TCSEC), whose cumbersome title was quickly reduced to "the Orange Book" from the color of its cover, became the center of the spectrum of "rainbow series" documents, each with its own distinct hue, that interpreted Orange Book requirements for various contexts (e.g., networks, databases) and provided guidance on where (in which risk environments) systems satisfying different levels of the criteria should properly be used. All of this was accomplished in an open environment.

The establishment of the NCSC also led to the consolidation of computer security research funding. There was natural concern among the military R&D funding organizations that the funds they had been applying to security R&D would be transferred to NSA, and they would lose control of them. Initially, services were allowed to retain control of their individual resources, but additional coordination was added to reduce duplication. Funds were also added to the NSA budget, leading to an overall increase in funds for generic computer security R&D. The process was relatively collegial, at least initially.

Did the military services actually reduce their investments in computer security research following the creation of the consolidated program? A more extensive study of the record than is possible here would be required to answer that question definitively. The perception that NSA had taken on the computer security problem could certainly have fueled competition in each service for those research funds that were being spent on computer security. A few years later, when some programs important to NSA were short of resources, NSA's own portion of the consolidated computer security program budget reportedly became the piggy bank. Nevertheless, throughout the 1980s, the resources available for computer security research generally seemed to be growing.

Several vendors bought into the Orange Book approach, developed systems, and submitted them for evaluation. At one point there were probably half a dozen or more efforts underway to develop different "secure Unix" implementations. Significant industry investment came into the area, fueled by expectations that the government would begin requiring systems meeting Orange Book criteria in its procurements. The demand for evaluators at the NCSC became hard to meet, particularly as vendors began hiring trained evaluators away from the government. In the meantime, other countries developed different models that moved the evaluation activities to independent laboratories. These laboratories would be self-financing by charging vendors for evaluating their products. NSA, while retaining a monopoly on evaluating relatively high-assurance systems, adopted this model as well.

My impression is that research funding for computer security declined in the early 1990s for several reasons. First, the end of the cold war brought general declines in military budgets and reduced funding for defense research. Second, the military embraced the idea that their computer systems could safely employ commercial off-the-shelf hardware and software, thereby avoiding the costs of development and maintenance for high assurance products. In effect, this decision implied that subversion – in the form of Trojan horses, trap doors, in-built logic bombs, and so on – was no longer considered a serious threat. So, the anticipated market for Orange Book products disappeared and industry naturally found it hard to justify further investment in them. Third, the outsourcing of evaluations reduced the demand for knowledgeable evaluators, whose training might involve participation in research projects. During the early 1990s, DARPA also seemed to decide that there were no great gains to be made in computer security and reduced their programs as well.

Again, relying on personal memory without documentation, my view is that Teresa Lunt's arrival at DARPA in the mid 1990s reflected the beginning of renewed funding (certainly from DARPA) in computer security research. By that time, intrusions were a common problem, and so intrusion detection and network security took on new prominence relative to host system hardening, When DARPA began to publish solicitations in the area of cybersecurity again, NSA took notice, leading to the creation of a joint program office between the two organizations. Subsequently, Teresa recruited Sami Saydjari from NSA to join DARPA in 1997; Sami in turn recruited Doug Maughan from NSA not long thereafter, and by 2000, DARPA had several sizable, unclassified information assurance research programs in progress. Sami estimates that DARPA invested a high of about $100M per year in defensive research by the year 2000. Along with Sami and Doug, DARPA PMs with programs in this area included Jay Lala, Cathy McCollum, Brian Witten, and Mike Skroch.

Although it had sponsored basic research in computer science and cryptography for many years, the National Science Foundation only started a focused program in cybersecurity in the fall of 2001, with the announcement of the Trusted Computing program, whose establishment was championed by Kamal Abdali and Helen Gill under Ružena Bajcsy. I was fortunate to be asked to get the program started with a budget of a few million dollars in 2001. But within a few years, partly in response to Congressional authorizing legislation (the Cyber Security R&D Act) under Greg Andrews, Wei Zhao, and Peter Freeman, I was managing the Cyber Trust program with a budget of more than $25M per year. In addition to four Cyber Trust center-scale awards, NSF funded the TRUST Science and Technology Center during this period, an award of about $5M per year for an initial five-year period. When my term at NSF ended in 2005, Karl Levitt assumed its management; the program has continued to grow under Ty Znati and Jeannette Wing's leadership. Last year it gained a new name, Trustworthy Computing program, and a renewed scope that adds emphases on privacy, usability, and foundations.

At the same time in 2001 that NSF decided to initiate its Trusted Computing program, DARPA acquired a new director, who allowed the existing information assurance programs to end and brought a focus on programs that could generate measurable, relatively near-term results. To many observers, the level of DARPA support for open research in the area appeared to decline substantially.

After completing his stint at DARPA, Doug Maughan joined the recently established Department of Homeland Security's Science and Technology division. He has continued to be a vital force in cybersecurity research funding, though DHS has focused on research aimed at producing results in the one to three year period.

IARPA and its forerunners, DTO and ARDA, also have invested funds in this area. Neither HSARPA, the DHS's version of DARPA, nor the newly created DoE ARPA-E, has shown significant interest in cybersecurity research.

The Comprehensive National Cyber Initiative, launched in 2007 under the coordination of Melissa Hathaway [4], focused strongly on near term measures to stem the tide of attacks against U.S. government systems and networks, but it also included an initiative for "leap ahead" research. This initiative has been responsible for some increase in funding to NSF (included in the CSIA figures) and is the source of funding for development of DARPA's planned National Cyber Range as well. As one of its first acts, the Obama administration called on Melissa Hathaway to lead a "60-day Cyberspace Policy Review" (followed by about 60 more days of coordination and revision prior to its public release) that focused on cybersecurity issues and the progress of the CNCI. The research community provided input to this review in several ways, including two white papers coordinated by NSF.

Largely missing from the foregoing narrative is the funding that has supported the DoD's internal laboratories (Naval Research Laboratory (NRL), Air Force Research Laboratory (AFRL), Army Research Laboratory (ARL)). Program managers Ralph Wachter at the Office of Naval Research (ONR), Bob Herklotz at the Air Force Office of Scientific Research (AFOSR), and Cliff Wang at the Army Research Office (ARO) have helped those organizations make well-considered investments in the area. Another significant source of relatively long-term research funds from the DoD in recent years has been the Multi-University Research Initiative (MURI) program. Steve King has provided leadership and coordination within the DoD's research funding complex.

## III. Concluding remarks

A 2002 report from the National Research Council was entitled "Cybersecurity Today and Tomorrow: Pay Now or Pay Later." It appears, that for a long time the US government wanted to wait until later to pay for needed research. The rising tempo of crime and serious, apparently state-sponsored, attacks in recent years seems to have raised its willingness to pay for some significant research efforts "now." Now that increased research resources do seem to be available, the research community must face up to two major questions: how to conduct research that can yield a significant, rather than incremental, improvement in the cybersecurity posture of critical infrastructures, and how to produce those results in such a way that they are easy for people to use. As was the case for automotive safety, however, real change to a safer cyber-infrastructure will likely require not only research and development but regulation in some form as well.

## Acknowledgments and Disclaimer

I thank Peter Neumann, Sean Peisert, Marv Schaefer, and Matt Bishop for inviting me to write this paper and for review comments that improved it. I also thank Sami Saydjari, Teresa Lunt, Steve Walker, Jeannette Wing, Ty Znati, Clark Weissman, Roger Schell, Paul Karger, and Fred Schneider for reviewing drafts, providing many helpful comments, details, and corrections, and for helpful conversations. I remain responsible for all opinions, errors, and omissions, and I apologize to the many people who have contributed to the funding of cybersecurity research over the years whom I have been unable to acknowledge here. The statements in this paper reflect my personal views only and not those of the University of Maryland or any part of the U.S. Government.

## References

[1] Figures taken from *Supplement to the President's Budget for Fiscal Year 2011*. A Report by the Subcommittee on National Information Technology Research and Development, Committee on Technology, National Science and Technology Council (and predecessor reports FY07 – FY10). Available at: http://www.nitrd.gov/pubs/bluebooks/index.aspx

[2] W. Ware, "Security Controls for Computer Systems: Report of Defense Science Board Task Force on Computer Security," Rand Report R609-1, February 1970. Available at

http://seclab.cs.ucdavis.edu/projects/history/papers/ware70.pdf

[3] J P. Anderson, "Computer security technology planning study," ESD-TR-73-51, vol. 1, ESD/AFSC, Hanscom AFB, Bedford, Mass., October 1972. Available at: http://seclab.cs.ucdavis.edu/projects/history/papers/ande72.pdf

[4] O. Sami Saydjari, "Launching into the Cyberspace Race: An Interview with Melissa E. Hathaway," *IEEE Security & Privacy Magazine*, Vol. 6, No. 6, (Nov-Dec. 2008), pp. 11-17.

Table 1. NITRD Cybersecurity Research Funding as Reported FY06 to FY10

| NITRD CSIA Budget supplement reports (Dollars in Millions) | FY06 Actual (estimate) | FY07 Request | FY07 Actual (estimate) | FY08 Request | FY08 Actual (estimate) | FY09 Request | FY09 Actual (estimate) | FY10 Request | FY10 Actual (estimate) | FY11 Request |
|---|---|---|---|---|---|---|---|---|---|---|
| NSF | 57.6 | 67.6 | 67.6 | 69.2 | 68.1 | 87.6 | 63.3 | 67.4 | 71.4 | 85.2 |
| DARPA | 78.7 | 81.6 | 93.4 | 96.9 | 124.4 | 106.8 | 125.4 | 143.6 | 143.5 | 126.1 |
| OSD and DoD Service research orgs. | 0.6 | 0.7 | 23.9 | 23.3 | 38.6 | 40.7 | 71.1 | 70 | 94.4 | 66.2 |
| NSA | 14.1 | 13.3 | 15.8 | 15.8 | 15.5 | 17.8 | 36.9 | 32.2 | 32.2 | 30.0 |
| NIST | 9.1 | 11.1 | 10.5 | 11.1 | 20.8 | 25.8 | 23.4 | 29.3 | 28.9 | 37.2 |
| NIH | 0 | 0 | 1.2 | 1.2 | 1.1 | 1.1 | 0 | 0 | 0.8 | 0.8 |
| NASA | 1.3 | 1.3 | 0.3 | 0.3 | 0 | 0 | 0 | 0 | 0 | 0 |
| DoE (Off. of Science) | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 3.5 | 3.5 |
| Totals | 161.4 | 175.6 | 212.7 | 217.8 | 268.5 | 279.8 | 320.1 | 342.5 | 374.6 | 349.0 |
| NIH – Recovery Act (ARRA) | | | | | | | 0.5 | | | |
| NSF - Recovery Act (ARRA) | | | | | | | 30.9 | | | |
| NIST - Recovery Act (ARRA) | | | | | | | 0.2 | | | |
| Totals with Recovery Act | | | | | | | 351.7 | 342.5 | | |
| CSIA was chartered in August 2005. First budget reports are FY06 actual/FY07 requests) | | | | | | | | | | |
| Unclassified research only. Other NITRD agencies report zero. | | | | | | | | | | |
| DHS, IARPA not included in NITRD agencies. | | | | | | | | | | |