# Looking Back: Addendum

David Elliott Bell

## Abstract

*The picture of computer and network security painted in my 2005 ACSAC paper was bleak. I learned at the conference that the situation is even bleaker than it seemed. We connect our most sensitive networks to less-secure networks using low-security products, creating high-value targets that are extremely vulnerable to sophisticated attack or subversion. Only systems of the highest security are sufficient to thwart such attacks and subversions. The environment for commercial security products can be made healthy again.*

## 1 Introduction

In the preparation of a recent ACSAC paper [1], I confirmed my opinion that computer and network security is in sad shape. I also made the editorial decision to soft-pedal criticism of my *alma mater*, the National Security Agency. At the conference, I found that I wasn't pessimistic enough and that there are additional disturbing developments.

Muting my criticism served neither my audience nor the national interest well.

## 2 Highlights of "Looking Back"

In today's networks, multilevel secure connections between networks are unavoidable.[1] The threat of sophisticated attack or subversion makes low-security systems unacceptable for connections between isolated networks. High-security systems must be used.[2] Government Off-The-Shelf (GOTS) products have never been satisfactory substitutes for commercial products. Before the Trusted Product Evaluation Program (TPEP), market forces alone never produced medium-security systems, much less high-security systems. A third way was needed — the changed business environment produced by Steve Walker's Computer Security Initiative.[3]

What we needed then and what we need now are "selfless acts of security" that lead to strong, secure commercial products.[4] Market-driven self-interest does not result in altruistic product decisions. Commercial and government acquisitions with tight deadlines are the wrong place to expect broad-reaching initiatives for the common good. Government must champion selfless acts of security separately from acquisitions.

## 3 NSA Has the Mission

In 1981, the Computer Security Evaluation Center was established at the National Security Agency (NSA). It was charged with publishing technical standards for evaluating trusted computer systems, with evaluating commercial and GOTS products, and with maintaining an Evaluated Products List (EPL) of products successfully evaluated. NSA was in effect appointed the Department of Defense (DoD) champion for high-security products and systems.

Changes in DoD and national policy have not fundamentally changed NSA's computer security mission (now, information assurance missions). NSA remains a government champion of standards and product evaluation of medium- and high-security systems.[5] [6]

## 4 Environment for High Security

Steve Walker's Computer Security Initiative intended to change the playing field for security systems. It succeeded. The combination of security standards, a list of evaluated

[1] This section is a précis of the relevant portions of my earlier paper. In the sections to follow, the supporting facts and reasoning have been relegated to appendices to simplify presentation. Appendix A is a plan of action. Appendix B is a summary of the results from last year's paper needed for this exposition.

[2] The author believes that the term "high assurance" has been debased by its frequent application to B1 and EAL4 systems, systems that provide no resistance to attack. In this paper, the term "high security" is used to refer to B3/A1 or EAL6/EAL7 systems. "Medium security" refers to B2/EAL5 systems. "Low security" refers to all other systems.

[3] Steve Walker, while at the Office of the Secretary of Defense, began the Computer Security Initiative to make major changes to computer security practice in the late 1970's. See appendix D.1 for more details.

[4] I use the term "selfless acts of security" to refer to actions or activities that have beneficial security results outside the immediate organization that performs them. Research into networking technology that led to the Arpanet and then to the Internet was a selfless act of networking. Developing a security metric and insisting on systems that measure up against that metric were selfless acts of security. Producing security archetypes or reference implementations of common networking components (file servers, dynamic web servers) would be selfless acts of security.

[5] The National Institute of Science and Technology (NIST) shares responsibility for running the National Information Assurance Partnership.

[6] See appendix C for NSA's original mission, policy changes since the Center's inception, and the current mission under DoD Directive 8500.1, "Information Assurance (IA)."

security products, and enforced policy to *use* evaluated products changed industry's calculus for decisions about "trusted" operating systems. IBM was disqualified from an NSA contract because it did not propose evaluated products. The DoD was serious about security and the marketplace took note.

Some of the archetype systems that had demonstrated and proved the consensus security principles were commercial products and underwent product evaluation (SCOMP at A1 and Multics at B2). Honeywell retained its Multics product line until it was cancelled in 1985. Honeywell (and subsequent owners of the technology) continued to use their own money to support and improve its SCOMP/XTS product line up to the current day. Private money was also put into developing new high-security systems. Gemini's Gemsos and Boeing's Secure LAN pursued A1 ratings successfully. Digital Equipment sought A1 for its Security Enhanced VMS (SEVMS) and was on a path to success when the product was withdrawn. Verdix, IBM Federal, and Trusted Information Systems pursued B2 ratings. Amdahl decided to submit its Multiple Domain Feature for evaluation at a B2 level, as did IBM with its PR/SM.[7]

With the transition from the *TCSEC* to the Common Criteria and with a shift in NSA's priorities to NSA-led security efforts like the Multilevel Information Systems Security Initiative (MISSI), the positive environment for high and medium security evaporated. The implicit assurance of swift accreditation for MISSI-based systems and an NSA imprimatur led to a rapid collapse of the commercial high-security marketplace. MISSI did not deliver on its promises and the business environment remains stunted.[8]

## 5 NIAP Performance

Since the inauguration of the National Information Assurance Partnership (NIAP), NSA has been the DoD lead for standards (the *Common Criteria* and protection profiles), new product evaluations, and maintenance of previously rated products. Its execution of standards work has been professional and thorough. New product evaluations have numbered one hundred twenty-seven, but only one of them was high security (EAL7) and another two were medium security (EAL5). The only evaluation against a validated U. S. Government Protection Profile was XTS-400's EAL5 augmented validation. The other validations were against product-specific Security Targets.[9]

NSA has not nurtured pre-existing high-security products well. The RAting Maintenance Program (RAMP) set up under the Trusted Product Evaluation Program (TPEP) provided a way to keep an evaluation rating of an evolving product current. RAMP was part of product evaluation — the phase immediately after the award of a rating. Since NIAP began, no effective method of transitioning existing high-security products into the NIAP family has been implemented. NSA/CCEVS's[10] position is that vendors should start over with a new *Common Criteria* evaluation. RAMP was intended to preclude complete re-evaluations of evaluated products. CCEVS's position is technically indefensible and reneges on the promise of RAMP.[11]

## 6 Competition with Vendors

Part of NSA's mission is to encourage high-security commercial products and their utilization. That encouragement has always been done through a combination of jawboning and leading by example. Leading by example includes both utilizing existing solutions in one's own operational systems and devising new solutions. The combination of urging the use of high-security components and using them yourself is a powerful argument to uncertain government players.

Unfortunately, there is an imbalance in NSA's leading by example. There has been much activity in devising new solutions under the Information Assurance Technical Framework (IATF)[12] and demonstration initiatives such as SELinux, NetTop, and recently the separation kernel.Not much has been evident in utilizing pre-existing high-security products for NSA projects. This is in contrast to the broad advocacy of NSA's products, most of which are low security. SELinux is not a trusted operating system according to NSA's SELinux FAQ web page[13] and commercial versions are a low-security EAL4. NetTop has not been evaluated but it builds on an EAL4 base. An EAL7 U. S. Government Protection Profile for a separation kernel is in development,[14] but its technical basis is unconvincing and frequently hard to follow.[15]

Since the technology transfer program for NSA's exemplar systems has been weak, NSA's internal initiatives are called "sandboxes" by critics, with some justification. Internal development activities should also address the need for high-security solutions that are evaluated at arm's length from the designers. They should then be effectively integrated into the marketplace to complement and not to compete with existing high-security commercial products.

[7] Neither Amdahl nor IBM received a B2 rating, but IBM received an EAL4 rating and several EAL5 ratings under the Common Criteria in Germany.

[8] See appendix D for additional details.

[9] One difficulty with the *Common Criteria* scheme is that product-specific Security Targets can be written to avoid hard security problems. As a result, even EAL7 does not assure one that a product does not have fatal flaws.

[10] Common Criteria Evaluation and Validation Scheme.

[11] See appendix E for details.

[12] "The IATF is a common reference guide for selecting and applying adequate and appropriate IA and IA-enabled technology in accordance with the architectural principles of defense-in-depth ..." [2]. DoDI 8500.2 directs NSA to "[m]aintain, update, and disseminate the Information Assurance Technical Framework (IATF) ... in coordination with the National Institute of Standards and Technology (NIST)" [3].

[13] Appendix F includes the relevant text from the FAQ.

[14] "U. S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness," dated 1 July 2004, is a protection profile "in development" [4].

[15] See appendices F, G, and H for details on SELinux, NetTop, and the separation kernel, respectively.

NSA's jawboning should be refocused on encouraging the use of commercial high-security products, rather than proselytizing for NSA's low-security demonstration systems.

## 7 Prevalence of Low Security

Very difficult computer- and network-security problems confront us today. This is best seen in the need for increased communication between different security domains and in the need to reduce desktop clutter caused by separate workstations for each different security domain. Both of these situations call for high levels of security, both from a policy perspective and from first principles.

Currently most of these critical needs are being filled by low-security components, whose deployments are contrary to DoD policy.[16] Even in the absence of similar policy, such deployments are ill-advised, verging on irresponsible. An attack on a weak connection no more sophisticated than the Cuckoo's Egg[17] gives an attacker complete control. Subversion of a weak connection allows adversaries to mount attacks of their choice at any time in the future. Such attacks can steal information, alter critical information, or bring down computers or entire networks. The worst subversion, called the "two-card loader," cannot be prevented or even found in low-security systems.[18] Direct attacks like the Cuckoo's Egg can only be thwarted by full mediation and self-protection, as required by B2/EAL5 or better systems. Complete prevention of subversion requires life-cycle security of the form found in A1 systems.[19]

There are many factors leading to this prevalence of low-security systems. One is the paucity of available high-security components. Another is a Morton's fork: either unawareness of these resources or a refusal to use them. Neither fork is justifiable.

NSA bears some responsibility for the low level of awareness because it has not jawboned for the right things and it has not led by example. DoD initiatives built on low-security platforms contribute to the problem.

Knowing about high-security resources and rejecting their use is nearly always wrong. On time, on budget, but unacceptably weak is not a defensible compromise. NSA could help this situation by crafting solutions for common problems utilizing existing high-security products.[20] Reference implementations of this sort would be a boon to program managers by lessening their technical, schedule, and certification risk. They could also be a springboard for updating the Information Assurance Technical Framework with content rather than the current placeholders for multilevel workstations, multilevel servers, and multilevel network components.

## 8 Conclusion

NSA's mission to nurture commercial high-security systems is languishing. At the same time, NSA resources are being expended in apparent, and sometimes explicit, competition with commercial vendors. These sandbox endeavors suffer from unacceptably low security. Existing high-security components are being overlooked or rejected in government initiatives where their use is crucial.

The United States should not be using low-security solutions in its most critical and most targeted components. The cost of a sophisticated attack or subversion here is extremely high. All players need to recognize the problem and take action to make sure that high-security components are used where they are needed.

NSA needs to re-focus on highly secure commercial products. It should incorporate existing high-security products into the NIAP family. It should encourage the development and deployment of new high-security products, especially through updates of the Information Assurance Technical Framework. NSA should redouble its efforts to encourage proper use, to lead by example, and to recreate the business environment where high-security products make good business sense.

All government organizations should utilize existing high-security products to solve their hardest security challenges: interconnecting single-level networks, providing multilevel secure workstations (both thin client and thick client), and providing timely information interchange among traditional enclaves of classified information. Relying instead on new solutions that (even if successful) will only be available in five to ten years, is not only impractical but also negligent of the public good.

## Acknowledgments

The insights of my readers made their peer review more like a collaboration. My thanks to Dave Amos, Jeremy Epstein, Paul Karger, Kenneth Olthoff, Tom Parenty, Ed Reed, Marv Shaefer, Roger Schell, Brian Snow, and Steve Walker.

[16] DoDI 8500.2 requires "controlled interfaces" for connections between a classified network and a lower-classified one. See appendices I, J and K for details on initiatives that are using low security rather than available high security.

[17] See appendix B.3 for a description.

[18] See appendix B.4 for details.

[19] The *Common Criteria* does not include strong requirements for trusted distribution. Instead, it has a single Delivery family, ALC_DEL, which requires that a developer document procedures for delivery and use those procedures. Hence, the EAL levels do not speak to trusted delivery at all. What one needs to prevent subversion of a high- or medium-security system is A1-style configuration management and trusted delivery.

[20] Such as data sharing between security domains, multilevel dynamic web servers, and multilevel thick and thin clients. See appendix A.2.

# A Plan of Action

This appendix is an abbreviated plan of action. The first subsection is an outline of the plan. Rationale for the course of action is provided in the second subsection.

## A.1 Summary

- Encourage commercial high-security products
  - Sponsor high-security reference implementations of common components
  - Establish an interim High-Assurance-Platform initiative using existing commercial products
  - Develop a high-security DoDIIS Trusted Workstation (DTW)
  - Develop a high-security Multiple Domain Dissemination System (MDDS)
- Advocate use of commercial high-security products
- Decrease advocacy of NSA's low-security initiatives
- Debate the technical merits of NSA's separation kernel
- Incorporate High-Security EPL Products into NIAP
  - Update or create new protection profiles to match B3 and A1
  - Vet the new or updated B3/A1 protection profiles
  - Add Gemsos, XTS-400, and Secure LAN to the Validated Products List.

## A.2 Reasoning

The positive business environment that emerged from the Computer Security Initiative built on four pillars:

- a standard way of specifying security requirements
- evaluation of products against those requirements
- policy to require use of evaluated products
- evidence that the policy would be enforced

The current situation is that there is a standard way of specifying security requirements (the validated U. S. Government Protection Profiles), a program of evaluating products (CCEVS) and policy to require use of validated products (DoDD 8500.1 and NSTISSP No. 11), but the evidence that the policy will be enforced is weak. This weakness shows itself in several ways:

- Advocating weak (EAL4) NSA products
- Advocating those NSA products over stronger commercial ones
- Designing new ten-year solutions but not dealing with interim needs

NSA should re-energize its role to encourage industry to produce high-security products as the DoD lead in NIAP, "promoting the development of technically sound security requirements for IT products . . . and appropriate measures for evaluating [them]" [5]. One useful initiative in this direction would be to assemble basic high-security components, built on existing high-security platforms. This effort should be simultaneous implementation on both Gemsos and XTS-400 with a cross-feed of lessons learned and shared source-code licenses for databases.

Two extremely common components are web-based connections between security domains and thin-client operations. For web-based connections, the order of implementation could be a multilevel file store, a static multilevel webserver, an extremely simple SeaViews-like multilevel database system, and finally a dynamic multilevel webserver. Multilevel thin-client architectures using high-security platforms can be achieved with sound security engineering. All the difficult problems for this initiative have already been worked out. What is required is to apply engineering talent and a modest amount of money. The experiences and the results could feed directly into updating placeholder sections in the IATF [6], as well as lead to high-security versions of the DTW and MDDS.

A second worthy initiative would be to address the same needs that the High Assurance Platform is addressing,[21] but sooner. A parallel effort to address these needs in the interim using commercial high-security products would be helpful to the user community at the same time as it reinforces NSA as the advocate for commercial high-security products.

A third initiative would be to increase the security of low-security DoD initiatives such as the Multiple Domain Dissemination System (MDDS) and the DoDIIS Trusted Workstation (DTW) by porting to or re-implementing on high-security commercial platforms.

A fourth initiative would be to establish a consensus concerning the separation kernel, or partitioning kernel, or "Multiple Independent Levels of Security" (MILS). Advocates should gather or produce publication-quality materials covering relevant technical topics such as definitions, descriptions, and rationale. Critics and advocates should then debate the issues publicly. After this trial by fire, either the concept will be tested and stronger or its flaws will be clear. In the interim, efforts to further proselytize or to finalize protection profiles should be put on hold.

As a last initiative, a review of existing U.S. Government Protection Profiles would be beneficial in assuring that all the important needs of the U.S. government are included as well as to determine whether there are updates needed. In this review, participants should include not only evaluators and Common Criteria staff but also vendors and village elders in the computer and network security community.

These initiatives would demonstrate NSA's commitment by publicly utilizing existing high-security products in vis-

[21] In a five- to ten-year timeframe.

ible endeavors and in leading by example. Combined with a scale-back in technical advocacy for NSA's low-security demonstration systems, the business atmosphere could be improved substantially.

# B Previous Results

## B.1 Intermediate Value Theorem

The Computer Security Intermediate Value Theorem (IVT) can be used to show that connections between single-level networks *are* multilevel.

**CS-IVT.** *If computer $A$ at security level $\alpha$ is connected to computer $B$ at security level $\beta$ through a network cloud and $\alpha \neq \beta$, then some processing platform in the cloud is multilevel.*

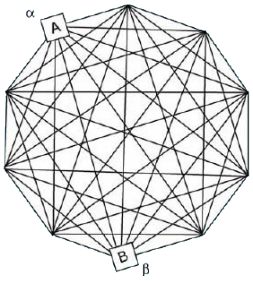

**Figure 1. Network Cloud**

***Proof Sketch*:** If it were not so, then each node on any path through the cloud to the distant end would be single-level at the same security level as its immediate predecessor. As a result, the first and last nodes would be operating at the same level, a contradiction [7].

In layman's language, every network path between two nodes of different security levels includes a multilevel node, strong or weak as the case may be.[22]

## B.2 Inevitability of Multilevel Connections

Consider two networks operating at two different levels, as shown in figure 2. When they are connected as shown, they become a single network. By the IVT, every path from the left network to the right network has a multilevel node on it. But every node in the individual networks is single-level by assumption. Thus, the multilevel node exists on the connection between the two networks.

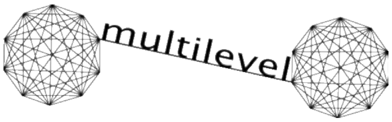

**Figure 2. Mostly Isolated Networks (MIN)**

For good reasons, the classified networks of the United States are currently operated as single-level networks with limited connections between them. Since those connections must be multilevel, one would prefer them to be as strong as we know how to make them. Unfortunately that is not the case.

## B.3 Cuckoo's Egg Attack

The Cuckoo's Egg attack described by Cliff Stoll was made possible by the ability of a user application to substitute the attacker's program for a system program (*atrun*) that ran every five minutes with superuser privileges. "The whole operation depended on his being able to move a file anywhere he wished" [8]. The initial step was to log on to an existing account with a stolen password. Logged on, the attacker created a bogus version of *atrun* then substituted it for the system's *atrun* using a feature of Emacs email. Unix actually allowed Emacs to replace another account's file without permission and without effective security checks. When the bogus *atrun* ran, it created a new superuser account for the attacker and restored the legitimate *atrun*. Success depended on a violation of presumed discretionary access controls: no one should be able to overwrite privileged code with code of their choosing. A properly configured, self-protecting B2 or better system would have prevented this substitution.

## B.4 Subversion

The Cuckoo's Egg breach was discovered by a smart geek who noticed a 75¢ discrepancy in the accounting. This kind of attack is bad, but a smart geek can find it. Subversions are much worse. In 1974, a subversion of Multics was demonstrated that required only ten 36-bit words [9]. Recently, the "two-card loader" has gained notoriety. The two-card loader is named after the mainframe loader that was punched into two cards (too large to fit on one card). The hardware reads the two loader-cards into a well-known location in memory, then transfers control to the first line of the loader program. The loader program, in turn, reads in the rest of the card deck and transfers control to the program contained therein.

A two-card loader subversion of an operating system reads in a malicious program as data, then transfers control to it. If a two-card loader is hidden in a commercial (or GOTS) operating system, then it can lie silently waiting for its trigger before doing its single, very simple job. A geek cannot find a well-written two-card loader. Exemplar subversions with a six-line "toehold" have been demonstrated [10], whereas Microsoft was unable to find an entire video game hidden in Excel before release.

[22] "Multilevel" here means processing information of more than one security level or classification at a time. This includes a node that is sequentially single-level, swapping out levels when they are dormant.

Systems meeting A1 requirements (specifically the configuration-management and trusted-distribution requirements) can prevent subversion through life-cycle control over design, source code, and the object code.

## C NSA's Mission

Steve Walker began the Computer Security Initiative in 1977 to pursue a third way to get truly secure computer systems into the marketplace [11]. The three parts of the initiative were a technical center for product evaluations, criteria against which evaluations could be performed, and a technical conference dedicated to computer-security issues. The first two invitational workshops were held in 1977 and 1978. Those workshops became the DoD/NBS Computer Security Conference.[23] In 1981, the Department of Defense Computer Security Evaluation Center was established. In 1983, the first version of the *Trusted Computer System Evaluation Criteria* (*TCSEC*) was published [12]. The Department of Defense version was published in 1985 [13]. The Trusted Product Evaluation Program (TPEP) was begun in 1981, using preliminary versions of the *TCSEC* which had been produced as part of the Computer Security Initiative.

### C.1 Original Direction

In 1982, Department of Defense Directive 5215.1 [14] gave the National Security Agency (NSA) the responsibility to create the Computer Security Evaluation Center "as a separate and unique entity within the NSA" and charged it to execute nine tasks, including:

1. "Establish and maintain technical standards and criteria for the evaluation of trusted computer systems . . . ."

2. "Conduct evaluations of selected industry and government-developed trusted computer systems against these criteria."

3. "Maintain and publish an EPL [Evaluated Products List] of the selected industry and government-developed trusted computer systems that is suitable for use by the DoD Components."

### C.2 Criteria

The *TCSEC* delineated seven classes of "trusted computer system," divided into four divisions. The best security that had been demonstrated was called "A1"; the other classes defined lower levels of security that could be used in situations not demanding the highest level of security. Security sufficient to thwart a frontal attack is found in systems from B2 up to A1.

### C.3 Evaluation Program

The Trusted Product Evaluation Program was a program to evaluate and record the fidelity of commercial products to classes within the *TCSEC* for later use by system certifiers and accreditors. Products successfully completing a "formal evaluation" were placed on the Evaluated Products List (EPL). TPEP and EPL constituted an institutional program of selfless security.

### C.4 Maintaining Ratings

One element missing from the original TPEP was the care and maintenance of successful evaluations. Since 1981 was at the beginning of the incredible shrinking product cycle, trusted systems faced a serious dilemma.

The dilemma lay in the need for trusted systems to be both commercially up to date and current in their evaluation. Since the product cycle was shrinking but the evaluation process was as short as it could be, trusted systems were behind the leading edge of commercial features and thus less appealing as building blocks. In addition, changes to a trusted system could require that at least part of its evaluation be re-visited. A complete re-write of the TCB would clearly have a substantial effect on the evaluation results. A change to non-TCB software would require no re-evaluation. Most changes fell between those extremes. Informed consideration of the changes, in the context of the security requirements and the previous evaluation, was required to determine the extent and nature of the re-evaluation effort.

TPEP dealt with maintaining evaluation ratings through the "RAtings Maintenance Program" (RAMP). A vendor employee (known as a Vendor Security Analyst or VSA) was vetted to monitor product changes in-house, with the ability to recognize and report both significant and minor security-related changes. The VSA periodically presented the changes to an NSA review board, which reviewed the RAMP materials. Depending on the review, NSA might accept the VSA's attestations as to the impact of the change and the testing performed, or could require additional government review.

When RAMP was under discussion, no one proposed having a vendor start over. The goal of TPEP was to put actually-secure systems into the marketplace and into the field. The evaluation process served that goal but did not supplant it. Since high-security systems are expensive to evaluate (tens of millions of dollars for an A1 system in the 1980's)[24] and take a long time to evaluate (ten years from initiation to final evaluation), the appropriate focus for RAMP was the potential for new security problems, in the context of the previous evaluation.

[23] Renamed the National Computer Security Conference and finally the National Information Systems Security Conference.

[24] NSA's figures for Gemsos's A1 rating were $14M for development and $50M for evaluation [15].

### C.5 Current Product Evaluation and Validation

With the transition from the *TCSEC* and TPEP to the *Common Criteria* [16] and NIAP,[25] NSA's mission stayed the same, but the details changed. Currently, responsibility for evaluating secure commercial products is split between commercial evaluation labs and NSA. The labs have responsibility for systems aspiring to low security (with assurance at EAL4 or below). NSA retains overall responsibility for medium- and high-security commercial products in the United States, at levels EAL5 through EAL7.

### C.6 Current Policy on Information Assurance

NSA's direction under current policy is substantially the same as it was originally.

1. "Establish and maintain technical standards and criteria for the evaluation of trusted computer systems . . . ."

   Under DoDD 8500.1, NSA is the DoD lead for NIAP, which includes responsibility for technical standards and criteria.

2. "Conduct evaluations of selected industry and government-developed trusted computer systems against these criteria."

   NSA validates evaluations under the *Common Criteria* against Protection Profiles and Security Targets.

3. "Maintain and publish an EPL of the selected industry and government-developed trusted computer systems that is suitable for use by the DoD Components."

   NSA maintains the Validated Products List (VPL), the modern equivalent of the EPL.

## D Business Environment for High Security

### D.1 Environment before TPEP

The problem that Steve Walker's Computer Security Initiative addressed was the failure of both methods of obtaining secure computer systems for DoD use — (a) using the security provided in commercial products and (b) crafting government-specific products, either as special versions of commercial products or as government-produced systems.

Commercial products did not provide any useful security. The "Tiger Teams" of the late 1960's demonstrated that attacking commercial systems always worked: "It is a commentary on contemporary systems that none of the known tiger team efforts has failed to date" [18]. Moreover, representatives of IBM, provider of 95% of the computers sold in the world, loudly and frequently asserted that technical security was bunk and that the only security required was personnel security.

Attempts to produce computers that met the DoD's needs for security had also proved ineffective. Paying a computer vendor to produce a special government-only version was a temporary fix. Fast-burner programmers wanted to work on new products, ones that would make their reputations [19]. Moreover, government-specific product versions remain available only as long as government funding continues. To keep up with new commercial releases, more government funds are required to update the government version. Internal government systems fared no better. Government does not employ system designers and implementors in the same numbers and at the same level of quality as do commercial companies. In addition, product support is not a natural government function and government does it badly.

One intention of the initiative was to change the environment for the conceptualization, design, and implementation of commercial computer systems. This was to be achieved by creating a standard metric for specifying security desired by the DoD and by promulgating policy that required the use of products that had measured up to the metric.

Before the initiative, government had scoped out the technical issues in computer security through full or partial funding of all the exemplar secure systems: Kernelized Secure Operating System (KSOS) [20], the Provably Secure Operating System (PSOS) [21], the Kernelized Virtual Machine (KVM) [22], and Multics [23]. These archetypes paved the way for business environment changes by demonstrating the technical feasibility of high-security systems.

### D.2 Environment under TPEP

With the *TCSEC*, the TPEP, the EPL, and prescriptive DoD policies, all the levers to alter the business environment for secure systems were in place. Multics and SCOMP (KSOS-6) were commercial products and they moved forward into formal product evaluation. DoD policy was issued, requiring all system acquisitions to have security requirements, phrased in terms of the *TCSEC* classes. Since policies are only effective if they are enforced, IBM's disqualification from NSA's Minstrel acquisition made the point. IBM began to bid other companies' products and all players took the policy seriously. As a result, the dynamics of business decisions changed where security was concerned.

Though profitable, Honeywell's B2 Multics stayed in the marketplace only till its cancellation in 1985. On the other hand, Honeywell continued to invest in SCOMP, moving from the original custom hardware through Level 6 Honeywell hardware, Intel 486/386 combinations, and Pentium, finally to Zeon. The software was continuously maintained and extended to include TCP/IP and middleware for the de-

[25] The National Information Assurance Partnership is a "joint initiative between the NSA and the National Institute of Standards and Technology responsible for security testing needs of both IT consumers and producers and promoting the development of technically sound security requirements for IT products and systems and appropriate measures for evaluating those products and systems" [17].

velopment of guard applications, all with internal funds.

Gemini Computers, Boeing, Verdix, IBM Federal, Trusted Information Systems, and Digital Equipment undertook new medium- to high-security products without government funding, based on their calculation that the products would have a market. Amdahl decided to undertake revisions of its Multiple Domain Feature (MDF) [24] sufficient to allow a product evaluation for the same reasons, hiring Trusted Information Systems to assist in the effort. IBM had developed PR/SM in competition with MDF. Later, it also pursued product evaluation [25].

NSA's actions in the 1980's, coupled with DoD and National policy, created an environment where secure products were viewed by the marketplace to confer competitive advantage in the ability to get government business. The marketplace responded with its own medium- to high-security products.

## D.3 Environment in Transition

In the early 1990's, an update to the *TCSEC* was begun. That effort went through several stages, resulting in harmonization of many nations' individual criteria into the *Common Criteria*. The result emphasizes product characterization over engineering methods and techniques proven through implementation and disputation. This untested revision to the community's security metric was a first ominous change.

A second indicator was the 1994 declaration by the Director of Central Intelligence and the Deputy Secretary of Defense that "Intelink [is] the strategic direction for Community product dissemination systems" [26]. Intelink used Internet technology to publish and read intelligence products. It joined independent agencies that lived under sometimes different policy regimes using low-security servers.[26] Not surprisingly, a prototype initially focused on allowing intelligence analysts to reference open resources such as Reuters and to exchange intelligence products was soon required to sanitize information and share it with the warfighter. In 1996, ASD/C4I asserted that "[e]fforts to provide more intelligence products for the warfighter at the collateral level must be a top priority" [28]. The security context for Intelink had changed dramatically. By then-existing (and current) DoD policy, those Intelink connections required high-security platforms. Low-security platforms, at best B1/EAL4, were put between intelligence and collateral networks. The prototyping and expansion of Intelink was a bad example in that the DoD was connecting separate security domains with low-security servers when high-security servers were both available and suitable for the task.

Then in 1995, the Multilevel Information Systems Security Initiative (MISSI) was announced. Its stated intention was to

> "be the security cornerstone for the defense information infrastructure. It will provide a multilevel security capability for networked automated information systems while ensuring that users can access only that information they are authorized, information is protected from unauthorized modification and users are identified and authenticated. MISSI will give us a single, integrated, consistent security infrastructure for all our needs: e-mail, electronic data interchange, electronic commerce, intelligence, and command and control, to include our business systems" [29].

Its "original goal was to provide a set of products and an architectural framework that would facilitate the development of multilevel secure NISs [Networked Information Systems]" [30].

Its effect was to kill the high-security marketplace.

MISSI was undertaken when the market boasted three evaluated A1 products, two B2 products, one promising A1 candidate, and two promising B2-like candidates. MISSI eventually reduced its goals to such an extent that its hallmark was the term "managed risk," a laudable-sounding concept that "has been used to justify use of low- or medium-assurance components to secure classified data (especially at the SECRET level) without much analysis of the threat or evaluation of the adequacy of the offered countermeasures" [31]. A1 vendors saw their business prospects collapse, virtually overnight. Boeing's Dan Schnackenberg believed that NSA through MISSI put all the Class A1 vendors out of business [32].

MISSI used the failed approach of funding government-specific products.[27] It also returned to an over-reliance on cryptography (Fortezza) over computer security.[28] It implied a quick path to accreditation: NSA was the official source of high-grade crypto *and* the executive agent for computer and information security. The promises of MISSI were never fulfilled and the changed environment created by the success of the Computer Security Initiative was destroyed. Those who really needed high security were the biggest losers when MISSI reneged and the high-security marketplace dwindled almost to nothing.

During the transition, the DoD and NSA in particular championed connections between different security domains, passing over available high-security products in favor of low-security ones. NSA then directly competed with high-security vendors with the introduction of MISSI, making promises that were not kept. Both by example and by explicit competition, the nurturing environment that grew out of the Computer Security Initiative was ruined.

[26] The required web statistics package Wusage 5.0 was ported to Solaris 2.4, SunOS 4.1.x, DEC's DigitalUnix 3.2, IBM's AIX, and Windows NT. AIX was B1 and Windows NT was C2. The others were unrated [27].

[27] Secure Computing Corporation's LOCK for the Secure Network Server and Trusted Mach for desktops.

[28] Computer security and cryptographic means supplement each other. Neither is a good substitute for the other. The best synergy results when they are used cooperatively and each is used to its strength.

### D.4 Environment Now

Today, there are three surviving medium- to high-security products, Aesec's Gemsos, Boeing's Secure LAN, and BAE Systems's XTS-400. They no longer have to compete with MISSI, but they do have to contend with the current NSA competition: SELinux, NetTop, NSA's separation kernel, and the High Assurance Platform.[29] In a sense, the situation is the same as it was in 1996. The positive business environment for high-security products and systems is gone and the player to beat is NSA. Experience shows that today's promises should be viewed with caution.

## E Performance under NIAP

Since the transition to the *Common Criteria*, NSA has had three responsibilities for trusted systems: (a) nurturing existing high-security products, (b) evaluating new medium- to high-security products at assurance levels from EAL5 through EAL7, and (c) participating in the Common Criteria working groups for the criteria and for processes and standards. Only the last task has received substantive attention.

### E.1 Nurturing Evaluated Systems

At this writing, there are only three product lines with *TCSEC* ratings of B3 or A1: Aesec's Gemsos (A1), Boeing's Secure LAN (A1), and BAE Systems's XTS-400 line (B3, a descendant of the A1 SCOMP).[30] The CCEVS organization[31] at NSA is responsible for this activity, but its current position on RAMP is that each system would have to re-establish its security *bona fides* by a complete re-evaluation under the Common Criteria.[32] This position flies in the face of common knowledge at the time that RAMP was established: only new security problems need be addressed. Insistence on procedure over substance does not serve the public interest. Since the Department of Defense now requires Common Criteria evaluations on products acquired, this lack of action on NSA's part has left the United States without the beneficial use of the best computer security products that have ever been produced.

### E.2 Evaluating New Systems

Since the establishment of the National Information Assurance Partnership (NIAP) in 1997, there have been 127 systems evaluated. Of those, only 3 were in the EAL5 to EAL7 range, and only one was against a validated U. S. Government Protection Profile.[33] One Tenix system (the Data Diode) was evaluated at EAL7 against a product-specific security target. A companion product from Tenix (the Interactive Link) was evaluated at the same time as EAL5 augmented, also against a product-specific security target. BAE Systems's XTS-400/Stop 6.1.E was evaluated at EAL5 augmented. In short, no new general purpose operating system[34] has been evaluated and listed since NIAP began.[35]

### E.3 Encouraging New Systems

NSA has not encouraged new commercial, high-security systems. Vendors have not spontaneously produced high-security systems. This is not a surprise, since as I discussed in my previous paper, market forces have never spontaneously produced high- or medium-security systems. NSA is more culpable than other players, since it has the responsibility to encourage new high-security technology and has not discharged its duties well.

### E.4 Common Criteria

Under NIAP, an international group is responsible for maintaining the *Common Criteria* itself and for identifying and resolving related problems. NSA and the National Institute of Science and Technology are the official United States representatives to this group and have executed their duties responsibly, with the possible exception of a new section on "composition" in the latest draft.[36]

## F SELinux

"Security-Enhanced Linux" (SELinux) was begun to demonstrate that labels and mandatory access control *could* be added to Linux code [35]. Later, type enforcement was added and the original purpose as an exemplar of one security feature (labels) was replaced with more general trusted operating system goals. Recent versions have a wide variety of mechanisms and settings, to the extent that an expert in SELinux internals is sometimes required to make changes to the configuration [36]. According to NSA's summary web page on SELinux,

> 12. Is Security-enhanced Linux a Trusted Operating System?
>
> No. The phrase "Trusted Operating System" generally refers to an operating system that provides sufficient support for multilevel security and evidence

[29] See appendices F, G, H, and I, respectively, for more details.

[30] Aesec, Boeing and BAE Systems are the current owners of the listed products. Only Boeing was the original recipient of the EPL rating.

[31] Common Criteria Evaluation and Validation Scheme.

[32] NSA's position on previously evaluated products was assembled from private comments by vendors, private comments by NSA representatives, and public statements by NSA representatives at the 2005 ACSAC conference.

[33] These figures are from [33].

[34] Nor M-component à la the *Trusted Network Interpretation* (*TNI*).

[35] IBM's evaluations of PR/SM were all performed in Germany [34].

[36] A quick reading of the draft text indicates that it stands in direct contradiction to my conclusions about composition requiring engineering constraints to succeed. If my reading is accurate, the draft text should be revised.

> of correctness to meet a particular set of government requirements. Security-enhanced Linux incorporates useful ideas from these systems but focuses upon mandatory access controls. This work is being combined with other efforts (e.g., auditing and documentation) to construct a "trusted" system that can be evaluated. The initial focus of Security-enhanced Linux development has been to create useful functionality that delivers tangible protection benefits in a wide range of real-world environments in order to demonstrate the technology [37].

SELinux is not listed on the Validated Products List, although Red Hat's Enterprise Linux is listed at EAL4. SELinux is nevertheless being touted and advocated for widespread adoption. Advocacy of SELinux by NSA staff directly competes with medium- and high-security products that can provide the same features and far greater assurance than does SELinux.

## G NetTop

NSA's NetTop [38] builds on SELinux. It is advocated as a high-assurance solution to replace multiple workstations on the desktop. Several vendors have NSA certificates and describe their NetTop products as EAL4, but no NetTop product is listed on the VPL.

Since NetTop attaches to networks of different security levels, DoDI 8500.2 would require it to be a "controlled interface," as "addressed in separate guidance" [39].

NetTop's reliance on SELinux's low level of security limits its assurance aspirations to EAL4, less than the requirement for a strong connection between different security domains.

The NSA imprimatur on weak-security NetTop products directly discourages the use of high-security workstations or thin clients attaching to domains at different security levels, a situation that DoDI 8500.2 says requires high security.

## H Separation Kernel

Since 1996, NSA staff have aggressively pressed an updated version of separation kernel under various rubrics — separation kernel, partitioning kernel and "Multiple Independent Levels of Security" or MILS.[37] A search of the open literature shows that little was published before 2002. Both the open and informal literature are replete with emphatic assertions but no justification, neither for the claims of prior failure nor for the optimism that This Time It Will Be Different. Add to that numerous factual errors in exposition and one can only wonder why NSA feels the topic is mature enough for codification as an official U. S. Government Protection Profile.[38]

The available and restricted literature do not provide a clear and concise definition of "separation kernel." It appears to be a small variation on Amdahl's MDF [44] and John Rushby's separation kernel [45]. Unlike those efforts, however, NSA's separation kernel has been tested neither in the marketplace (like MDF) nor in the marketplace of ideas (like Rushby). Assertions that high security can be achieved without trust (in the *TCSEC* sense) have been strongly rebutted. Paul Karger's thesis, similar to Rushby's separation kernel work but less widely known, makes the point clearly that multilevel components are required to realize the concept.[39]

Advocates of NSA's separation kernel assert that mutilevel security is not needed; data isolation, control of information flow, periods processing and fault isolation suffice.[40] This position is undermined by [48] which points out that since all boolean lattices are isomorphic (up to the number of atoms), all boolean policies are similarly isomorphic. The traditional classification/category policy is a boolean policy. Moreover, any policy that can be described using uninterpreted symbols and *AND*, *OR*, and *NOT* is also a boolean policy. Isolation can be so described. Isolation is a boolean policy. Thus, isolation is isomorphic to traditional non-discretionary (mandatory) access control. In summary, isolation is not different from nor simpler than multilevel security.

NSA's separation kernel and the related MILS have gotten far ahead of themselves. There is no clear and concise definition. There is little open literature on it. What literature exists is too much hortatory and not enough analytical. It does not build on community archetypes. It has not been subjected to critical peer review. In spite of these hurdles unjumped, setting it in concrete – whatever "it" is – is eighteen months along.

Success in getting adherents on this bandwagon cannot have rested solely on its technical basis.[41] It has probably benefited from hard problems in avionics and real-time operating systems that the separation kernel, if successful, would solve. It has also benefited from NSA's position and reputation.

Similar promises about separation kernels in the past have come to little, even with the efforts of such as John Rushby. It would be best for NSA and for users of the separation kernel to get all the ducks in a row *before* institu-

[37] In [40] published in 2004, the work is described as on-going for ten years. A 2005 paper from the University of Idaho [41] references "High Assurance Architecture via Separation Kernel" [42], an informal communication dated 1996.

[38] According to the NIAP website, the "U. S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness" [43] is currently "in development" as a U. S. Government validated Protection Profile.

[39] "... the network communications processors are multilevel secure ...." [46].

[40] See [47] for example.

[41] It is worth noting that enthusiasm for the separation kernel is not universal in the embedded and real-time communities. Pointed criticism of the draft "Partitioning Kernel Protection Profile" includes the following: (1) "PKPR ignores real-time"; (2) "PKPR passes over most of the Common Criteria functional requirements"; (3) "PKPR reliance on a simple use of memory protection hardware mechanisms rules out most of the RTOS market"; and (4) "PKPR requires that validating each component provides for compositional correctness." Regarding (4), "This is simply incorrect in any but the simplest systems" [49].

tionalizing it as a protection profile and *before* widespread advocacy for its adoption.

Adherents believe this separation kernel to be a great new hope for security. Let them document their case in a form suitable for publication. Let advocates and critics have a debate on the issues. After a trial of fire, the separation kernel will emerge hardened and improved; or its flaws will be manifest.

## I High Assurance Platform

According to its flyer, "[t]he High Assurance Platform (HAP) Program is a National Security Agency (NSA) initiative whose main objective is to develop the technologies and standards that provide a highly assured computing environment for applications" [50]. The primary capabilities are to provide access to multiple domains from a single workstation and to allow secure data movement between domains. According to DoDI 8500.2, this kind of connection requires a controlled interface [51].

Although HAP focuses on workstations and servers, its *purpose* is to enable connections of these platforms to multiple security domains and to transmit data between them. Their approach is multilevel hypervisors, a topic of some subtlety [52]. In essence, HAP intends to insert high-assurance code (the assured information sharing service) between the hypervisor (untrusted) and applications (untrusted) [53]. This kind of appliqué, inserting a security layer between untrusted parts, has never worked and certainly is not the basis of actually high security. Experience has shown that evolving COTS C1/D class materials into a PL5-accreditable system is an impossibility.

There are at least two problems with the current HAP program. One is the lack of appreciation of unconstrained composition.[42] HAP needs global network conditions and a network worldview as a framework into which individual platforms can fit.

The second difficulty is the rejection of existing high-security systems as solutions to this problem.[43] This is particularly troublesome, since normal development and deployment cycles mean that current HAP efforts will not be fielded for at least five – and more likely ten – years.

Lacking a *TNI*-like perspective, HAP still believes that it can solve unconstrained composition when no one else has succeeded. I believe it may be impossible. The laudable goals of HAP are undercut by the refusal to use proven technologies in the short term, preferring untested technologies only available five to ten years from now. Its approach conflicts with DoD policy about connections between security levels.

## J Multiple Domain Dissemination System

The Multiple Domain Dissemination System (MDDS) is a multilevel server that provides cross-domain interactions via the user's web browser [55]. Its principal component is a product called WebShield. WebShield is hosted on Trusted Solaris, an EAL4 operating system. WebShield is designed to attach to two or more networks simultaneously and to pass `http` requests made in one domain to other domains, returning the response to the originator. Searches can be cloned to all relevant security domains and the results collated and displayed.

Deployment of WebShield between networks of different classifications would be disallowed by current DoD policy [56]. EAL4 systems are vulnerable to Cuckoo's Egg attacks and two-card loader subversion. They are not strong enough for use in MDDS. MDDS's deployment puts at risk all networks to which it connects.

## K DTW

According to the Rome Lab web site, the DoDIIS Trusted Workstation is

> "a certified multi-level security (MLS) information system utilizing both commercial and government off-the-shelf technologies. Composed of server hardware running a trusted operating system and ultra thin-client devices at each user's desktop, the DTW provides simultaneous access to multiple security domains from a single workstation. DTW additionally provides a secure means of transferring data between security domains" [57].

The "trusted operating system" mentioned is the EAL4 Trusted Solaris. Current DoD policy would disallow DTW's connection of its Citrix servers to multiple security domains [58]. Connecting DTW's weak-security servers to two or more security domains is technically weak and introduces severe vulnerabilities, not only for direct users but also for users of the attached networks.

## References

[1] D. E. Bell, "Looking Back at the Bell-La Padula Model", *Proc.* ACSAC, 7–9 December 2005, 337–351.

[2] DoD Instruction 8500.2, "Information Assurance (IA) Implementation," February 6, 2003, ¶E3.2.4, p. 32.

[3] *Ibid.*, ¶5.6.6, p. 5.

[4] U. S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness," version 0.621, NSA/IAD, 1 July 2004.

[5] DoD Directive 8500.1, "Information Assurance (IA)," October 24, 2002, p. 22.

[42] Two A1 systems *could* be combined to make a A1 combination. Improper combination would make the result less than A1. Artful failure of requirements could result in any *TCSEC* class between D to A1. Composition without global engineering constraints may not be possible. See [54]
.

[43] At the 14 August, 2006 HAP Vendors Forum, HAP management made it clear in its opening remarks that references to or advocacy for existing high-security products were not welcome and that such products would not be considered within HAP.

[6] "Information Assurance Technical Framework (IATF)," Release 3.1, September, 2002, ¶6.7.2, ¶6.7.3, and ¶6.7.4, pp. 6.7-23.

[7] Bell, *op. cit.*, p. 312.

[8] C. Stoll. **The Cuckoo's Egg**. Doubleday: New York, NY, 1989, p. 29.

[9] P. A. Karger, R. R. Schell, "Multics Security Evaluation: Vulnerability Analysis," ESD–TR–74–193, Vol. II, ESD/AFSC, Hanscom AFB, MA, June 1974.

[10] C. E. Irvine, "Considering Lifecycle Subversion," OSD Invitational MLS Workshop, Alexandria, VA, 24 September, 2003.

[11] S. T. Walker, "The Advent of Trusted Computer Operating Systems," *Proc.* National Computer Conference, May, 1980, 655–665.

[12] Department of Defense Trusted Computer System Evaluation Criteria, CSC-STD-001-83, 15 August 1983.

[13] Department of Defense Trusted Computer System Evaluation Criteria, DOD 5200.28-STD, December 1985.

[14] Department of Defense Directive 5215.1, "Computer Security Evaluation Center," October 25, 1982, change 1, November 16, 1994.

[15] R. R. Schell, "High Assurance MLS Systems through Proven Technology," AFCEA MLS Panel, Omaha, NB, 24 May 2005.

[16] *Common Criteria for Information Technology Security Evaluation*, Version 3.1, CCMB-2006-09-001, September, 2006.

[17] DoD Directive 8500.1, *op. cit.*

[18] J. P. Anderson, "Computer Security Technology Planning Study," ESD–TR–73–51, Vol. I, AD–758 206, ESD/AFSC, Hanscom AFB, MA, October 1972, p. 4.

[19] T. Kidder, **Soul of a New Machine**. Little, Brown: Boston, 1981.

[20] E. J. McCauley and P. J. Drongowski, "KSOS—The design of a secure operating system," *Proc.* AFIPS 1979 NCC **48**, 345–353.

[21] P. G. Neumann, R. S. Boyer, R. J. Feiertag, K. N. Levitt, and L. Robinson, "A provably secure operating system: The system, its applications, and proofs," Technical Report CSL–116, SRI International, 1980.

[22] M. Schaefer, R. R. Linde, *et al.*, "Program Confinement in KVM/370," *Proc.* ACM National Conference, Seattle, October, 1977.

[23] J. Saltzer, "Protection and the Control of Information in Multics," *Comm. ACM* **17**(7), July 1974, 388–402.

[24] R. W. Doran, "Amdahl Multiple-Domain Architecture," *Computer* **21**(10), October, 1988, 20–28.

[25] P. A. Karger, private communication, 2006.

[26] R. J. Woolsey (DCI), J. M. Deutch, (DepSecDef), Memorandum for Co-Chairmen, Intelligence Systems Board, Subject: Intelink, 11 August 1994.

[27] F. T. Brown. **Top Secret Intranet.** Prentice Hall: New Jersey, 1998.

[28] E. Paige Jr., "The Rapid Expansion of Intelink," **Defense Issues**, **11**(66), June 11, 1996.

[29] E. Paige Jr., "From the Cold War to the Global Information Age," **Defense Issues**, **10**(34), February 27, 1995.

[30] F. Schneider, *ed.*, **Trust in Cyberspace.** National Academy Press: Washington, DC, 1998, box 4.4.

[31] *Ibid.*, p. 109.

[32] D. Schnackenberg, OSD Invitational MLS Workshop, Alexandria, VA, 24 September, 2003.

[33] A. Dale, "National Information Assurance Partnership (NIAP) Common Criteria Evaluation and Validation Scheme (CCEVS)," Briefing to IAPB, 23 March 2006.

[34] "PR/SM$^{TM}$ LPAR for the IBM System z9$^{TM}$ Enterprise Class and the IBM System z9$^{TM}$ Business Class from International Business Machine Corporation (IBM)," Bundesamt für Sicherheit in der Informationstechnik, BSI-DSZ-CC-0378-2006, September, 2006.

[35] NSA's SELinux FAQ, NSA web site, 2006.

[36] "An introduction to SELinux," LWN.net, Article 103203, 2006, with comments.

[37] NSA's SELinux FAQ, *op. cit.*

[38] NetTop Technology Profile Fact Sheet, NSA web page, 2006.

[39] DoDI 8500.2, *op. cit.*, E4.A4, Enclave and Computing Environment, p. 85.

[40] R. Beckwith, W. M. Vanfleet, L. MacLaren, "High Assurance Security/Safety for Deeply Embedded, Real-Time Systems," Embedded Systems Conference, 2004.

[41] J. Alves-Foss, C. Taylor, P. Orman, "A Multi-layered Approach to Security in High Assurance Systems," *Proc.* 37$^{th}$ Hawaii International Conf. on System Sciences, 2004.

[42] P. White, M. Vanfleet, C. Dailey, "High Assurance Architecture via Separation Kernel," internal communications, 1996.

[43] NSA/IAD, "U. S. Government Protection Profiles for Separation Kernels in Environments Requiring High Robustness," version 0.621, NSA, Ft. Meade, MD, 1 July 2004.

[44] Doran, *op. cit.*

[45] J. Rushby, "A trusted computing base for embedded systems," *Proc.*, 7th DoD/NBS Computer Security Initiative Conference, Gaithersburg, MD, September 1984, 294–311.

[46] P. A. Karger, "Non-Discretionary Access Control for Decentralized Computing Systems," M.I.T. Laboratory for Computer Science, MIT/LCS/TR-179, May 1977, p. 54.

[47] W. M. Vanfleet *et al.*, "MILS: Architecture for High-Assurance Embedded Computing," STCS CrossTalk, August, 2005.

[48] D. E. Bell, "Lattices, Policies, and Implementations," *Proc.* 13$^{th}$ National Computer Security Conference, Washington, DC, 1–4 October 1990, 165–171.

[49] V. Yodaiken, "Reponse to 'Partitioning Kernel Protection Profile,' " unpublished draft.
`[http://www.yodaiken.com/papers/securityx.pdf]`

[50] HAP Flyer, August, 2006.

[51] DoDI 8500.2, *op. cit.*

[52] P. A. Karger, "Multi-Level Security Requirements for Hypervisors," *Proc.* ACSAC, 7–9 December 2005, 267–275.

[53] HAP flyer, *op. cit.*

[54] Bell, "Looking Back at the Bell-La Padula Model," *op. cit.,* ¶6.2.

[55] W. Neugent, "Assured Information Sharing," *The Edge,* The MITRE Corporation, Fall 2005, **9**(2), 12.

[56] DoDI 8500.2, *op. cit.*

[57] Rome Lab web site: `www.rl.af.mil/tech/programs/jedi/`

[58] DoDI 8500.2, *op. cit.*