DEPARTMENT OF THE AIR FORCE
UNITED STATES OF AMERICA
MCMXLVII

AIR UNIVERSITY review

JANUARY-FEBRUARY 1979

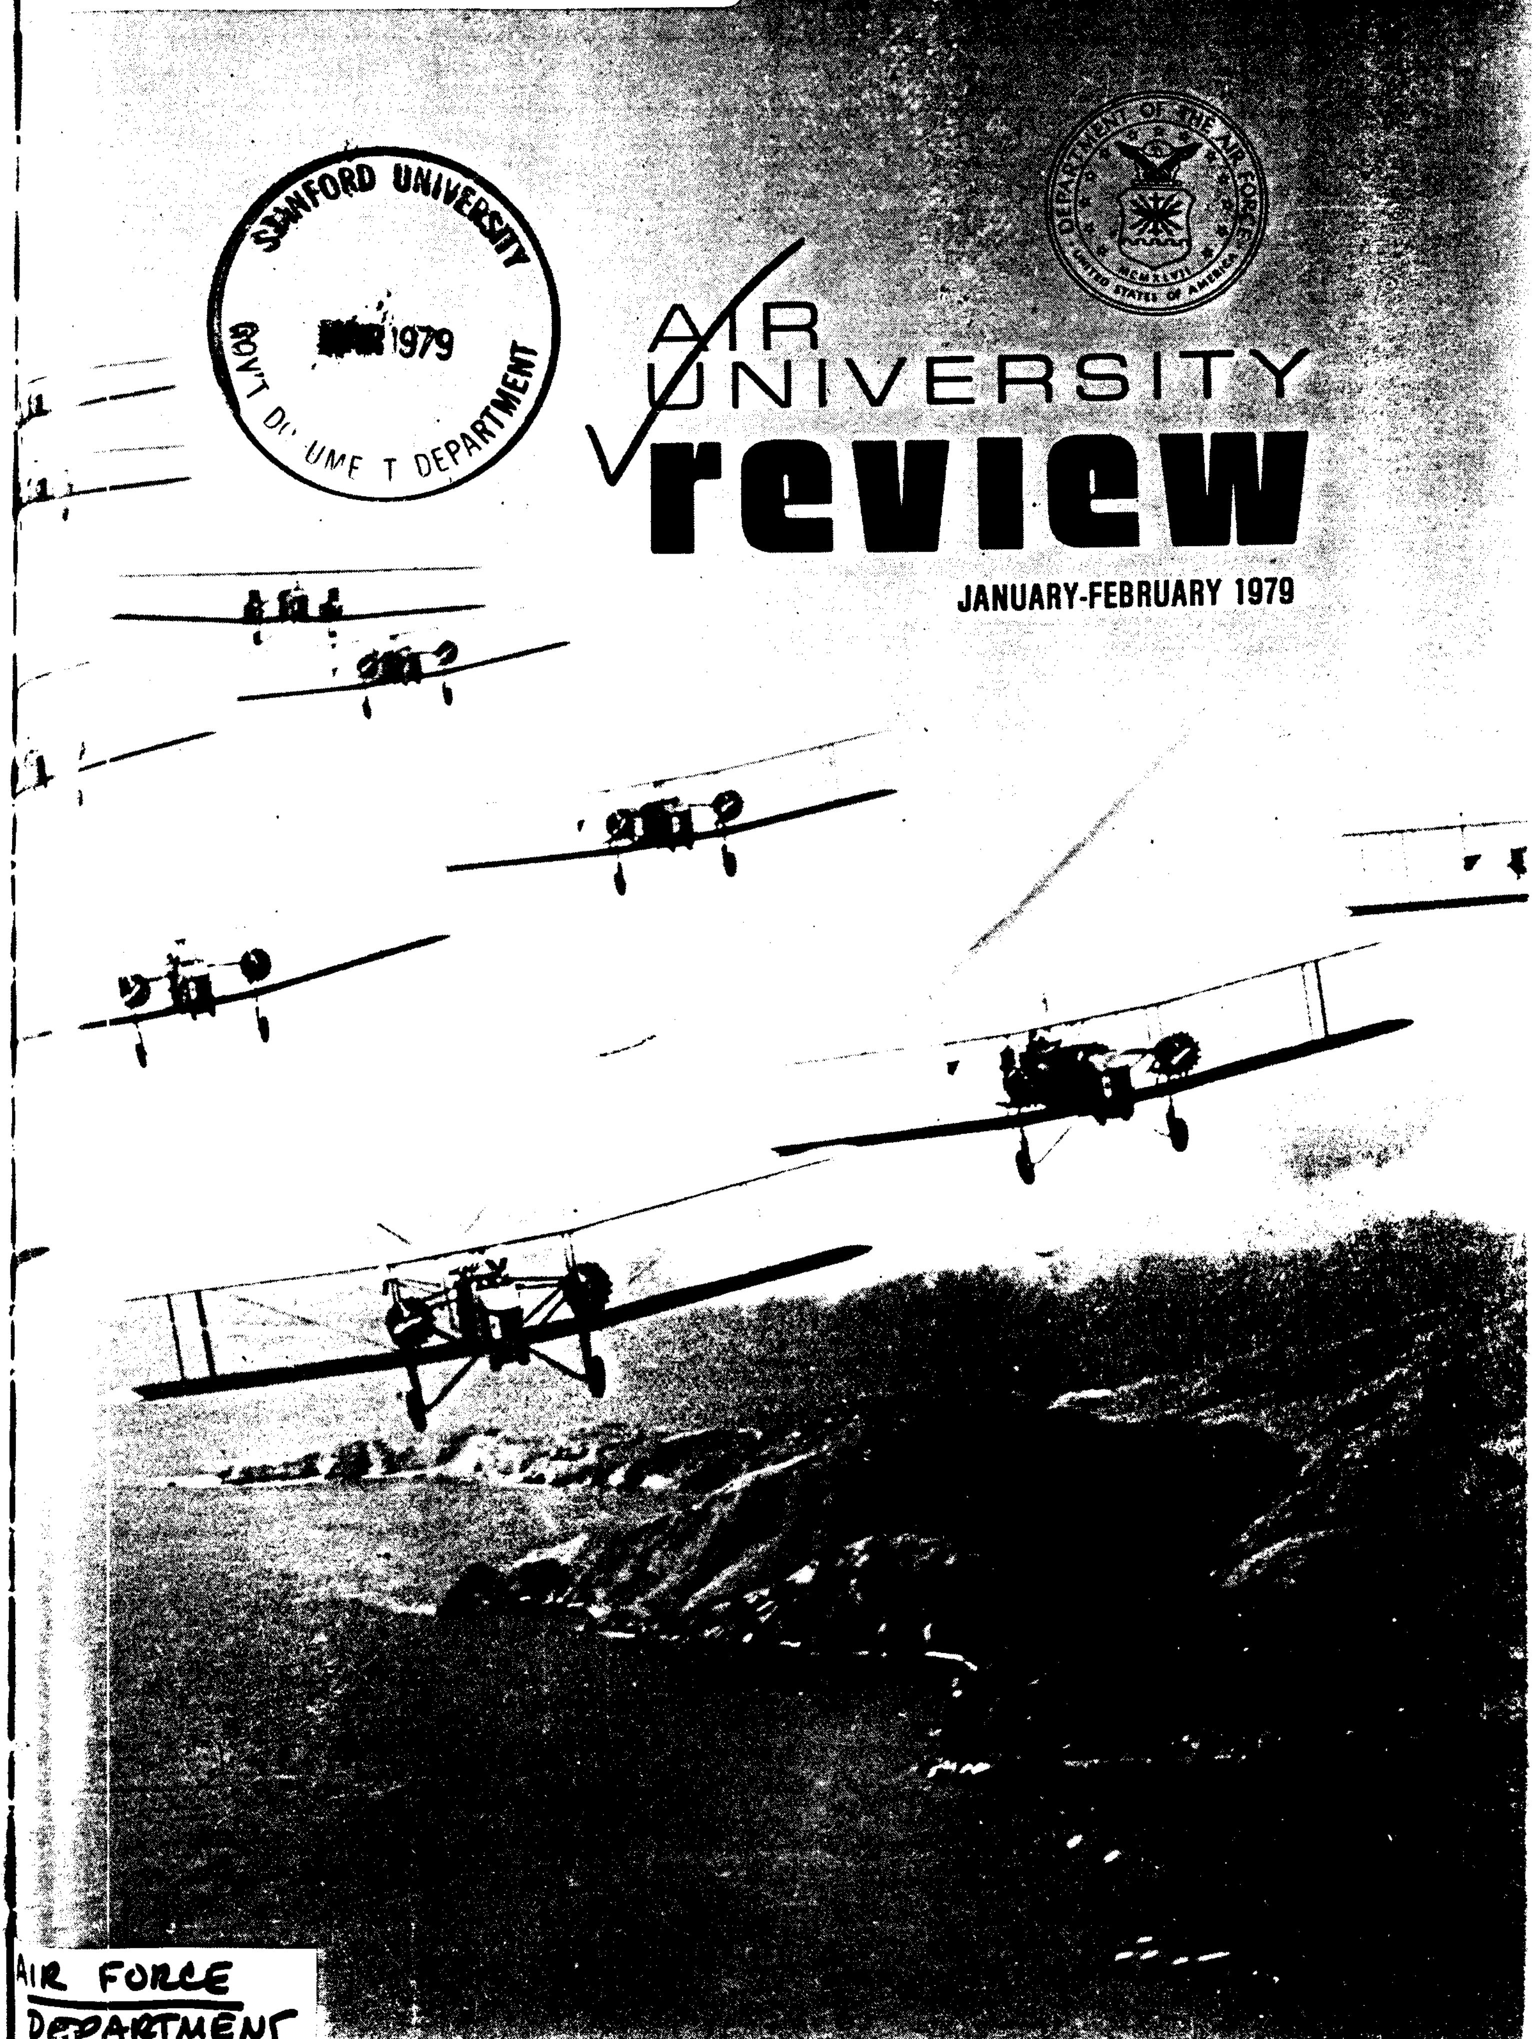

COMPUTER SECURITY

the Achilles' heel of the electronic Air Force?

LIEUTENANT COLONEL ROGER R. SCHELL

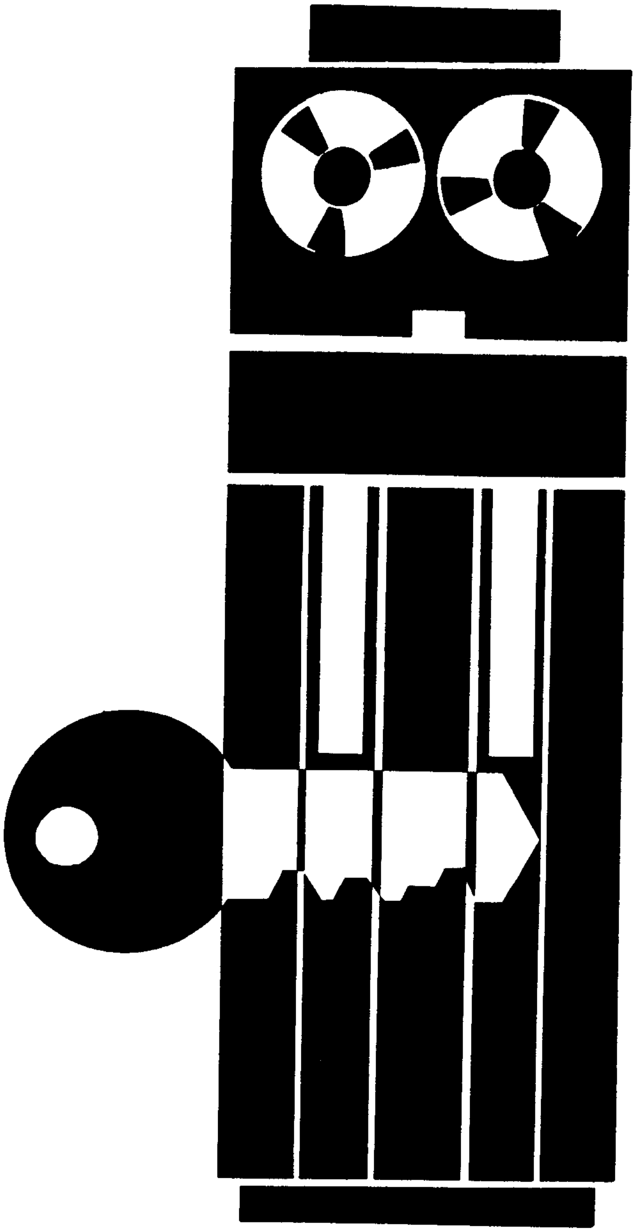

THE KGB officer addressed the select group of Soviet officials with his usual tone of secrecy but an unusual air of excitement:

> Comrades, today I will brief you on the most significant breakthrough in intelligence collection since the "breaking" of the "unbreakable" Japanese and German cyphers in World War II—the penetration of the security of American computers. There is virtually (if not literally) no major American national defense secret which is not stored on a computer somewhere. At the same time, there are few (if any) computers in their national defense system which are not accessible, in theory if not yet in fact, to our prying. Better still, we don't even have to wait for them to send the particular information we want so we can intercept it; we can request and get specific material of interest to us, with virtually no risk to our agents.
>
> The Americans have developed a "security kernel" technology for solving their problem, but we need not be concerned—they recently discontinued work on this technology. They are aware of the potential for a computer security problem, but with their usual carelessness they have decided not to correct the problem until they have verified examples of our active exploitation. We, of course, must not let them find these examples.

Your first reaction to this scenario may be, "Preposterous!" But before you reject it out of hand, recognize that we know it could happen. The question is: Will we apply sound technology and policy before it does happen? To be sure, there are things we do not know about the probability of success of such an effort, but we can rationally assess the most salient controlling factors:

- The high *vulnerability* of contemporary computers has been clearly indicated in the author's experience with undetected penetration of security mechanisms. In addition, security weaknesses are documented in both military and civil reports.
- The *capability* of the Soviets (or any other major hostile group) to accomplish the required penetration is quite evident. In fact, no particular skills beyond those of normally competent computer professionals are required.
- The *motivation* for such an information collection activity is apparent in prima facie evidence. The broad scope and high intensity of Soviet intelligence efforts in areas such as communication interception are frequently reported.
- The potential *damage* from penetration is growing with the ever increasing concentration of sensitive information in computers and the interconnection of these computers into large networks. Through computer penetration an enemy could, for example, compromise plans for employment of tactical fighters or compromise operational plans and targeting for nuclear missiles.
- The *opportunity* for hostile exploitation of these vulnerabilities is increasing markedly both because of the increased use of computers and the lack of a meaningful security policy controlling their use. In the name of efficiency many more people with less (or no) clearance are permitted easier access to classified computer systems.

We have a problem and a solution in hand. Detailed examination of a hostile nation's (e.g., Soviet) capability and motivation in those areas is properly in the realm of the intelligence analyst and largely outside the scope of this article. However, it will trace the outlines of the computer security problem and show how the security kernel approach meets the requirements for a workable solution—although recent termination has nipped in the bud very promising work toward a solution.

What Makes Computers a Security Problem?

Although a certain appreciation of subtlety is needed to understand the details of the computer security problem, our objective here is to illuminate the basic underlying issues. To understand these issues, I will examine not only the capabilities and limitations of computers themselves but also their uses.

First, we take for granted the fundamental need to protect properly classified sensitive military information from compromise. Security has long been recognized as one of the basic principles of war, and throughout history security or its lack has been a major factor of the outcome of battles and wars. We can and do strictly control information when the dissemination is on paper. It is, therefore, illogical to ignore the fact that computers may disseminate the same information to anyone who knows how to ask for it, completely bypassing the expensive controls we place on paper circulation.

Second, we must appreciate that "exploitation of the phenomenal growth of computer science is a major area of technological emphasis within DoD."[1] We currently lack quantitative superiority (or even parity) in several force level areas, and computers appear to be able to provide the qualitative superiority we must have.

The need for these capabilities is clear when we realize that "good C^3 [command, control, and communications] capabilities can double or triple force effectiveness; conversely, ineffective C^3 is certain to jeopardize or deny the objective sought."[2] Indeed, we have in a very real sense become an "electronic Air Force"[3] with computers at our heart.

Finally, we need to recognize that some major vulnerabilities may accompany the substantial benefits of computer technology. Most decision-makers cannot afford the time to maintain a thorough understanding of explosively developing computer technology. But they can even less afford to be ignorant of what the computer can do and also of how it can fail. In particular, a commander responsible for security must ensure that dissemination controls are extended to computers. He must be able to ask proper questions—to surface potential vulnerability for critical and unbiased examination.

historical lessons in emerging technology

It is not new to find that an emerging technology is a mixed blessing. In particular, the threat facing computers today is illustrated in the evolution of military electrical communications—an earlier revolutionary technology. Our compromise of the security of Axis communications was fundamental to the outcome of World War II, and computers now offer our enemies the opportunity to turn the tables on us.

Military communication specialists early recognized the vulnerability of electrical transmission to interception, e.g., through wire taps or surreptitious listening to radio signals. The solutions were simple and effective but drastic: restrict transmission only to relatively unimportant (viz., unclassified) information or to transmission paths physically guarded and protected from intrusion. Likewise, for several years the Air Force restricted computer use to either unclassified data or to a protected computer dedicated to authorized (cleared) users. In both instances the security solutions limited use of the technology where most needed: for important information in potentially hostile situations, such as battlefield support.

The communication security restrictions gave rise to various cryptographic devices. These devices were to encode information into an unintelligible and thus unclassified form so that protection of the entire transmission path was not required. But (of paramount importance to us here) this dramatically changed the very nature of the security problem itself: from a question of physical protection to a question of technical efficacy. The effectiveness of the cryptographic devices was argued, based not on careful technical analysis but rather on the apparent absence of a known way to counter them. Presently, computer technology is in a position analogous with a similar argument for its effectiveness against unauthorized access to computerized information. In both instances, the arguments seem to offer an acceptable risk in spite of a de facto weak technical foundation.

Technically weak cryptographic devices found widespread military use because of false confidence and the pressing operational need for electrical communications. One notable example was the Enigma machine used by the Germans during World War II. Their high-level national command and control network used it for communication security throughout the war. As *The Ultra Secret* records, "the Germans considered that their cypher was completely safe."[4] Yet, before the war really got started, the British had in fact "solved the puzzle of Enigma."[5] The Air

Force is developing a similar dependency with each (formal or de facto) decision to accredit computer security controls. In either case policy decisions permit a technical weakness to become a military vulnerability.

Examples during World War II show how the tendency to defend previous decisions (to accept and use mere plausible techniques) assures the enemy of opportunities for exploitation. In Europe the broken Enigma signals (called Ultra) "not only gave the full strength and disposition of the enemy, it showed that the Allied [troops] could achieve tactical surprise."[6] In fact, General Dwight Eisenhower stated that "Ultra was decisive."[7] *The Codebreakers* describes a similar misplaced trust by the Japanese and notes that American cryptanalysts "contributed enormously to the defeat of the enemy, greatly shortened the war, and saved many thousands of lives."[8] To be sure, the Germans "must have been puzzled by our knowledge of their U-boat positions, but luckily they did not accept the fact that we had broken Enigma."[9] Similarly, the Japanese "hypnotized themselves into the delusion that their codes were never seriously compromised."[10] The Axis establishment, it seems, would not acknowledge its security weakness without direct confirming counterintelligence—and this came only after they had lost the war. As for Air Force computer security, the absence of war has precluded ultimate exploitation; yet, the lack of hard counterintelligence on exploitation has already been offered as evidence of effective security.

Although technical efforts led to these devastating vulnerabilities, it was nonetheless the technical experts like William Friedman who provided a sound technical basis: "His theoretical studies, which revolutionized the science, were matched by his actual solutions, which astounded it [the scientific community]."[11] Today our military makes widespread use of cryptographic devices with confidence. For computers, as for communications, the nub of the problem is the effectiveness of the security mechanism. Recent logically rigorous work has resulted in a security kernel technology. However, DOD is not yet applying this technology.

The thrust of this historical review is captured in the maxim, "Those who cannot remember the past are condemned to repeat it." The historical parallels are summarized in Table I. The main lesson to be learned is this: Do not trust security to technology unless that technology is demonstrably trustworthy, and the absence of demonstrated compromise is absolutely *not* a demonstration of security.

distinction between computation and protection

A given computer in one installation may securely handle sensitive data, and an identical machine may be totally insecure in another installation. The key to understanding the computer security problem is to distinguish when the computer provides only computation and when it must also provide security. These are two very distinct cases.

In the first case, commonly called "dedicated mode," the computer and all its users are within a single security perimeter established by guards, dogs, fences, etc. By the use of secure communications, this perimeter may be geographically extended to remote terminals. Only these external security controls are required to maintain the security of the system. Use of the computer is restricted so that at any time all the users, remote or local, are authorized access to all the computerized information. A potential attacker must overcome the external controls and penetrate the inner sanctum of cleared per-

	Electrical Communications	Electronic Computers
Limited Use	unclassified only protected paths	unclassified only dedicated facility
Plausible Security	cryptographic technology crucial to security no known counter weak technical foundation	internal security controls crucial no known penetration weak technical foundation
Unwarranted Dependence	false confidence in cryptography policy acceptance	false confidence in internal controls policy acceptance
Underestimated Enemy	repeated and undetected interception advocates demand counterintelligence	repeated, undetected, and selective access advocates demand counterintelligence
Adequate Technology	information theory	security kernel

Table I. Comparative evolution of security problems

sonnel. The computer provides only computation; no failure or subversion of the computer itself can compromise security because of the protected environment.

In the second case, commonly called "multilevel mode," the computer itself must internally distinguish multiple levels of information sensitivity and user authorization. In particular, the computer must protect some information from certain users. For multilevel mode, internal security controls of hardware and computer programs must assure that each user may access only authorized information. For multilevel security the computer itself must clearly provide protection as well as computation. For the potential attacker, simply gaining access to the peripheral users of the computer will suffice—if he can penetrate the internal controls.

Multilevel security controls function analogously to a cryptographic device; their effectiveness is central to information security. Because of the inherent structure of computers, a multilevel security weakness invites repeated exploitation. Furthermore, those security failures internal to the computer are almost certain to be undetected. In contrast to communications where enemy access to important traffic is a matter of chance, in a penetrated computer he has selective access, not only for extraction but also for modification of information of his choosing. All the worse, the processing power of modern computers provides this information rapidly and completely.

If we are worried about protecting our cryptographic codes, then we are indeed foolish to neglect our computers. And we must realize that multilevel mode can aid the attacker unless the internal controls of the computer itself provide reliable protection.

evidence of weak security controls

The critical question then is this: Dare we trust the internal security controls of computer programs and hardware? The author's experience with security weaknesses indicates that contemporary computers do not provide reliable protection. Computers proposed as sufficiently secure to protect sensitive information were checked for security shortcomings. A formally sanctioned "tiger team" looked for weaknesses in these supposedly secure computers. (For accuracy the examples will be limited to those evaluations in which the author personally participated.)

The tiger team operated as a legitimate user with only limited access to a small part of the information in the system. The team objective was to penetrate internal security controls and demonstrate that unauthorized access could be gained. In every instance of the author's experience, serious security weaknesses were discovered after only a few hours or days of effort.

Passwords for the asking. A common element of protection is a secret password or key that the user must provide in order to receive services or information. To be effective the secrecy of the passwords must be preserved. An IBM 370 computer with the time-sharing option (TSO) had remote terminals in various uncontrolled areas; the secret passwords restricted the users' access. This particular computer contained sensitive Air Force procurement source-selection information with tightly controlled dissemination. The tiger team members found that they had merely to ask by name for the password file and the passwords for all the TSO users would be printed for them—without a trace that the passwords had been compromised. The designers had overlooked the relationship between security and the ability to print a file.

Good commercials not enough. In the Pentagon a General Electric system called "GCOS" provided classified (secret) computation for the Air Staff and others with secured remote terminals at selected locations. The manufacturer made an advertising thrust about his security. Air Force advocates proposed making a multilevel system by adding unsecured remote terminals, for unclassified uses, for better coordination and efficiency. Again, passwords were to protect the sensitive information. When a user presented his password to the computer, GCOS checked a list of passwords to verify the user's legitimacy. To make this check, GCOS copied part of the list into its main memory. Among other flaws, the tiger team found that GCOS left this copy of the passwords where it could be printed easily and without trace. The designers had overlooked the possibility of deliberate misuse of a necessary computer function.

Government designers not perfect. After the Pentagon penetration, some advocates claimed that government designers with a greater awareness of security could avoid such flaws. An organization that processed sensitive intelligence data spent a substantial effort "fixing" basically the same GCOS system. They were confident they could maintain multilevel mode security. The tiger team found that these "fixes" could easily be circumvented. In this case not only could any user get at any information in the system but also he could access the classified information in computers connected in a network with that computer!

A contract cannot provide security. Basically the same GCOS system was selected for a major command and control system. Advocates assured the users that it would be made multilevel secure because security was required by the contract. An extensive tiger team evaluation found there were many deep and complex security flaws that defied practical repair—the computer was finally deemed not only insecure but insecurable.

The best security is not good enough. Honeywell Information Systems, with DOD sponsorship, modified the GCOS computer in an effort to improve several areas substantially, including security. The resulting Multiplexed Information and Computing Service (Multics) was widely touted for its security. The tiger team used an Air Force laboratory computer to evaluate Multics as a potential multilevel secure computer for the Pentagon. Although it had the best security design of any system encountered, the tiger team found several implementation flaws.[12] In one case Multics first checked a prospective user's authorization for access to information and, when the request proved valid, executed the request. However, the user could change the request after the validity check but before execution; Multics then executed the changed request, allowing unauthorized access. This penetration of Multics came from an implementation short cut made to improve efficiency.

Encrypted passwords retrieved. The Multics system internally encrypted its password list so that even if printed out the passwords were not intelligible. When a user presented his password, it was encrypted and then compared to the encrypted list. The tiger team used the penetration technique developed on the laboratory computer to access the encrypted password list of a large university and then broke the cypher to obtain all the passwords.

Trap door installed. The tiger team penetrated Multics and modified the manufacturer's master copy of the Multics operating system itself by installing a trap door: computer instructions to deliberately bypass the normal security checks and thus ensure penetration even after the initial flaw was fixed. This trap door was small (fewer than 10 instructions out of 100,000) and required a password for use. The manufacturer could not find it, even when he knew it existed and how it worked. Furthermore, since the trap door was inserted in the master copy of the operating system programs, the manufacturer automatically distributed this trap door to all Multics installations.

Audit record destroyed. Some have argued that a computer need not always prevent unauthorized access as long as it keeps an audit record of such accesses. The Multics system kept a protected audit record of access, and the tiger team's unauthorized accesses were recorded. However, the audit record was itself subject to unauthorized access. The tiger team merely modified the record to delete all trace of its actions, such as insertion of the trap door.

Even fixes have holes. Honeywell produced a new Multics computer that corrected all the implementation flaws reported by the tiger team. The tiger team used Honeywell's new computer at their Phoenix, Arizona, manufacturing plant and penetrated the security again.[13] This new flaw resulted from changes made to correct the previous ones! It was becoming increasingly clear that providing a multilevel secure computer was indeed difficult.

Trojan horse not dead. While some had recognized the problem, advocates in the Air Staff were commending an installation for their multilevel security solution on another computer. The solution consisted of programs to segregate the classified and unclassified information. There were no remote terminals, but users could submit unclassified jobs to the computer without security checks. From an unclassified job the tiger team penetrated the underlying computer operating system and modified the solution into a Trojan horse, an apparently useful program that concealed harmful capabilities. The Trojan horse hid an invisible copy of classified jobs. A later unclassified job retrieved the hidden information, compromising security. Thus the security solution was not only ineffective but it actually exacerbated the security problem.

The obvious moral. Few if any contemporary computer security controls have prevented a tiger team from easily accessing any information sought. These examples are by no means exhaustive; they must not be used to infer predominance of certain flaws or to associate particular weaknesses with only a few manufacturers. Others have comparable security problems.

futility of evaluation by penetration

In a very real sense the Air Force has been fortunate that security is so poor in current computers—the greater danger will come when the argument that a computer is secure because tiger teams failed to penetrate it appears plausible. Indeed, evaluating internal computer security controls is a most difficult challenge. As with cryptography, there are basically two approaches.

If the security controls are based on a carefully formulated, sound technology, then they may be subject to rational analysis of their effectiveness. As already noted, this is generally not true of contemporary computers. The security kernel approach, which is subject to such methodical technical analysis, will also be discussed.

Alternatively, an advocate can simply search for ways to penetrate a computer's controls; failing to penetrate, he can plausibly argue there is no way to penetrate since none is known (to him). If a security hole is found, it can first be patched before arguing for security. Obviously, this argument suffers acutely from both theoretical and practical difficulties.

In principle, one could test all possible programs to find any that led to a security penetration. This method of exhaustion would be effective but is far beyond the realm of feasibility. For any substantial computer this would take so long that before the evaluation was finished the sun would literally have burned out! Thus, a realizable evaluation by exhaustion must be so incomplete as to be ludicrous.

In fact the effort spent in penetrating and patching yields poor marginal return in terms of security. The tiger team examples indicate some of the difficulties:

First, experience shows that new penetrators tend to find new holes—even after previous teams have found all they could. It seems unlikely that a real attacker will not involve new people.

Second, holes do not generally result from rank stupidity but from human oversight in dealing with a difficult design problem. Thus the fixes themselves are likely to be flawed.

Third, it does not take a highly specialized expert to penetrate security. It is true that most computer professionals do not know ways to penetrate the systems they use; they want to do a job, not interfere with it. Yet when given the assignment, even junior and inexperienced professionals have consistently succeeded in penetration.

Fourth, the exposure to attack is frequently much greater than from just the known system users. Commercial telephone connections to military systems are increasing and give worldwide access. Communication taps also give access to unsecured direct connections; microwave intercepts by the Soviets in the U.S., as recently revealed by the White House, demonstrate this capability. Lack of strict security control on the submission of computer jobs allows attacks in the name of a legitimate user even for computers without remote terminals. Interconnection to other computers can add a large group of unknown users as well.

Fifth, the attacks can be developed and perfected on other than the target computer. A similar computer owned or legitimately accessed by the attacker can be used to minimize the risk of detection. Once perfected, the attack methods can be applied to the target computer.

Finally, to a hostile penetrator the trap door and Trojan horse approaches are probably the most attractive, and these deliberately created flaws in computer programs are the most difficult to detect. Most tiger teams concentrate on accidental flaws that anyone might happen to find, but the deliberate flaws are dormant until activated by an attacker. These errors can be placed virtually anywhere and are carefully designed to escape detection. Yet most military systems include programs not developed in a secure environment, and some are even developed abroad. In fact some systems can be subverted by an anonymous remote technician with no legitimate role in the system development. These errors can be activated by essentially any external interface—from an unclassified telegram to a unique situation set up for detection by a surveillance system.

ON BALANCE, penetrating and patching internal controls is not a promising security technique. Even without the prospect of trap doors and Trojan horses and without military security demands, "private companies have attempted to patch holes in so-called [secure] computer systems, and after millions of dollars and years of effort, they gave up in failure."[14] This approach is little more than a game of wits in which the designer must try to find (and patch) *all* the holes while the enemy need find (and exploit) but one remaining hole—a rather unbalanced contest.

The "bottom line" is simple. The commander responsible for security in a computer system needs an unequivocal answer to one crucial question: Is security dependent on internal controls? That is, is there any failure or subversion of the computer itself that could degrade security? If so, with contemporary computers he has a root inconsistency in the laxity about computer security within the military environment that normally has strict controls on dissemination of sensitive information.

Computer Security Alternatives

We have seen that in contemporary computers the internal controls are not only ineffective but also defy assessment. Yet obviously we can choose to follow the path of the German and Japanese cryptographic experience—underestimating enemy exploitation of the technical weaknesses. This is the chance we have taken in each of several Air Force decisions to operate contemporary computers in a multilevel mode.

If we lose this gamble, the damage depends on what the computer is protecting. It can range from violation of personal privacy to fraud, battlefield damage, or pre-emptive surprise attack. For example, it has been proposed that the Air Force dynamically retarget its strategic ballistic missiles; this supports the national policy of flexible response and would allow application of retaliatory weapons to the most lucrative military targets. However, computers are at the heart of this capability; if they were penetrated, an enemy could retarget the missiles to impact on low-value or even friendly targets as part of a surprise attack!

We will not attempt to explore the numerous possible scenarios from dependence on weak techniques, but we will look at solution alternatives. Both technical and policy issues are involved. Basically, the Air Force has two alternatives other than to ignore the problem: either limit computer use or use available adequate technology to make the internal controls reliable.

avoid dependence on internal controls

The obvious alternative is to deliberately restrict computer use to a dedicated mode so that the internal controls cannot affect security. There are three common ways to avoid dependence on internal controls.

First, a separate computer can be dedicated to each level of classified information. This is particularly attractive for an on-line or real-time system where the information must be immediately accessible. This approach can lead to duplicate or inefficiently used computers.

Second, each level of classified information can be scheduled to use the computer for a different time period. This requires purging of information from all the system memory at the end of a scheduled period. This usually cumbersome manual procedure lacks responsiveness and wastes computer resources while the change in classification level is completed.

Third, various classification levels can be processed together. All communication lines must be protected, and all the users would need to be authorized access to all the information. Since the internal controls are not dependable, all output from the system is tentatively classified at the highest level. For information with a lower classification, a competent authority must manually review the output for contamination and downgrade it before releasing it at the lower level.

These use restrictions can support good security, but they result in a substantial degradation of capability in a modern computer.

Added expense. These security restrictions significantly add to the cost. Additional communication security measures are needed, and additional manpower is required for the manual review of output. There is also the cost of security clearance investigations for the users whose information the computer may contaminate with information of a higher classification. Other costs include those for duplicate equipment and for additional capacity to compensate for wasted resources. For example, when one major computer system failed to deliver the promised multilevel security, major Air Force sites had to clear many users and make multimillion dollar purchases of additional equipment.

Increased risk. In practice the dedicated mode leads to a major increase in the exposure of information. The lack of internal controls effectively destroys the compartmentalization intended to limit the damage from subversion. The greater number of people requiring clearance increases the chance of granting access to an untrustworthy individual. Manual purge procedures are prone to errors that leave classified memory residues which can be extracted by unauthorized users. Furthermore, the manual review of large volumes of computer output may in fact be a bureaucratic ruse to transfer security responsibility (liability) from designers to users; the reviewer has little chance of detecting unauthorized classified information that has been accidentally or deliberately included in the output.

Foregone capabilities. Such security restrictions can seriously limit the operational capability of battlefield support systems. Modern weapons demand command and control systems with rapid access to a large base of current and accurate information. This (necessarily shared and integrated) data base will typically contain information ranging from unclassified through top secret. Since many people who maintain the less classified information have limited clearances, and the volume of information requires that computers be used, we have the classical multilevel computer security problem. Internal computer controls are crucial to information protection, and avoiding dependence on the internal controls will seriously limit system capabilities.

The problem is exacerbated by interoperability with its interconnected network of computers with a large, diverse, and geographically dispersed user community. Command and control system computer networks are a prime example. Yet one military official observed that because of poor internal computer security in one such network, its 35 large-scale, general-purpose computers would never truly be used for the purpose for which they were procured. The problem is even further intensified by the growing need for fusion of selected intelligence information (without compromise of sensitive sources) with tactical operations information.

In summary, the dedicated mode avoids many computer security problems but fails to meet the operational needs of a modern military force. These needs can only be met by effective multilevel protection in the computer itself.

apply adequate technology

Developing and applying reliable internal computer security are neither easy nor impossible. Although the need for multilevel operation is frequently recognized, the military has given only limited attention to developing the required technology. In fact, the Air Force recently directed termination of its multilevel security development program, the largest in the Department of Defense.[15]

Before we examine the technological progress that has been made, it should be instructive to identify some of the reasoning that surfaced in the recent Air Force termination. The pattern of thought reflects that computer security is not currently a major focus.

- The prospect of industry's solving the computer security problem is overestimated by concluding that industry has the same security problem as the military. However, the communications analogy indicates a difficulty. In the civilian sector, communication security violations are subject to legislation, not prevention; wiretapping is outlawed, and there is legal redress for loss. In contrast, the military must resort to prevention (e.g., military approved cryptography), since we cannot sue the KGB! The computer situation is similar; there are legislative thrusts but limited commercial success toward demonstrably effective internal controls. The wait for spontaneous industry solutions is likely to be a long one, and it is unlikely that they will ever meet military security standards in areas such as protection from deliberate subversion.

- Inadequate research and development (R&D) funding was allocated to continue one element of the program at an optimal level. Yet portions of the program with funds available were also terminated. Eight million dollars of work was successfully completed. About $10 million of work over four years remained to complete development of a full prototype and the associated general basis for competitive procurement. Several estimates indicate that development costs could be recouped by avoiding the penalties of dedicated mode—not to mention the increased security and operational capability.

- The threat is minimized by seeking counterintelligence that is practically unavailable, e.g., actual examples of enemy agents caught in the act. The enemy may appear too ignorant for penetration, not interested in military secrets, or incapable of planned subversion and exploitation. A single number quantification of the probability of threat can implicitly assume a random incident rather than a planned penetration activity. This may indicate acceptable risk without an objective criterion of acceptability. These perceptions are generally not based on professional intelligence methods with “worked examples” (e.g., from communication security) of the methodology.

- Interest in developing solutions is limited by a lack of clear responsibility for the effectiveness of internal controls. Staff and policy offices can provide recommendations, guidance, and even approvals for computer security mechanisms without responsibility (liability) for any security compromise that might result. On the other hand, the security test and evaluation efforts and cost-effectiveness assessments of individual commanders are largely unrelated to the system's real protection. This is in marked contrast to military communication security where technical experts are responsible for certifying the security mechanisms.

- The computer security problem is difficult to recognize when policy does not clearly distinguish the cases where the computer simply provides computation and where the computer provides internal protection. Such policy focuses development on security controls that are “not necessarily certifiably perfect”—a rather ambiguous goal. In such a policy framework requirements analysis will not identify the need for internal controls. In fact, a computer may well satisfy all regulations and still be highly vulnerable.

- Confidence in weak controls grows from the assumption that expending resources on security will substantially improve security. In fact, the effort may be simply ineffective, as in the case of the penetrate and patch treadmill. Current policy enumerates computer design characteristics for internal security that are neither necessary nor sufficient for security.

- Attention to security gimmicks results in overlooking serious weaknesses. There are many mechanisms of minimal effectiveness in improving internal security controls—handprint analyzers, encryption of internal data, read-only memory for security information, etc. Some guidance has encouraged computer programs that sort out and label products by security level. Evaluation of these programs focuses on expected results with friendly users rather than on deliberate subversion of the programs or penetration of the underlying system. Pursuing such scattered efforts is frequently worse than doing nothing at all, since it gives a dangerous false sense of security.

THESE SORTS of issues caused the Air Force to characterize its Electronic Systems Division's (recently terminated) development program as “controversial.” But our previous examination of the problem makes it clear that multilevel operation without adequate technology is a high stakes gamble. Most charitably, it is strangely inconsistent with established standards in other areas (e.g., communications) of military security that hypothesize a deliberate, competent, and motivated hostile threat and respond with effective countermeasures. More likely it nullifies all other security measures, allowing damage limited only by the imagination of the enemy.

Security Kernel Technology

Fortunately, military R&D—in particular the recently terminated Air Force program,[16]—has made substantial progress toward adequate technology for multilevel security. A major step toward solution was the introduction in 1972 of the security kernel[17] technology, which provided a scientific foundation for demonstrably effective internal security controls. Although an explanation of the technical details is well beyond the scope of this article, one technical report summarizes the kernel approach this way:

> The approach to obtaining a secure system involves first defining the security requirements and then creating a conceptual design that can be shown to provide the required protection (i.e., a model). The model formally defines an ideal system (in our case one that complies with military security requirements), and provides a basis for testing a subsequent implementation. Once a [security kernel] that meets the requirements previously described has been implemented, computer security has been achieved. Of the software in the system, only the security kernel . . . need be correct. . . .The operating system proper and/or the application software can contain inadvertently introduced bugs or maliciously planted trap doors without compromising security.[18]

Under the Air Force program the security kernel demonstrated its technical feasibility, independent of any particular computer vendor or security policy. The kernel has also largely established its operational acceptability, with specific evidence for broad functionality, good efficiency, security certifiability, and supportability. In addition, the underlying technical requirements of the kernel have been successfully incorporated into military procurement specifications for both a commercial large-scale computer and an embedded weapon system computer. In short, the basic technology is well in hand.

scientific foundation

A security kernel is a small set of computer program instructions and associated hardware that controls all access by users (viz., through their programs) to information. A given security kernel is usually unique to a particular computer. A security kernel for computers is in many ways conceptually analogous to a cryptographic device for communications.

Security kernel design is derived directly from a precise specification (viz., a mathematical model) of its function. (The kernel model is analogous to the algorithm that defines the mathematical function of a cryptographic device.) This mathematical model is a precise formulation of access rules based on user attributes (clearance, need to know) and information attributes (classification). System parameters control an installation's specific use (e.g., for the DOD classification policy, privacy protection, etc.).

The chief distinguishing characteristic (from whence its name) of the security kernel concept is that a kernel represents a distinct internal security perimeter. In particular, that portion of the system responsible for maintaining internal security is reduced from essentially the entire computer to principally the kernel. Thus the kernel is analogous to a cryptographic device that removes most of a communication path from security consideration. To be a bit more technical and concrete, a typical security kernel has several (say ten to twenty) small computer programs (viz., subroutines) that can be invoked by other programs (e.g., the operating system and individual user application programs). The kernel, and only the kernel, controls and manages all the hardware components that store and access information. All other (viz., nonkernel) programs must invoke the kernel (i.e., call on its subroutines) in order to access information—the kernel checks the user and information attributes and provides only access that is authorized. Yet, in spite of these checks, there is minimal user impact. Figure 1 conceptually illustrates this structure.

The technical breakthrough was the discovery of a set of model functions and conditions that are provably sufficient to prevent compromise for all possible nonkernel computer programs. Each function of the model determines the design for a kernel program. In addition, the model imposes security conditions that must be met by the design. Security theorems have been proved showing that (since the kernel precisely follows the model) the kernel will not permit a compromise, regardless of what program uses it or how it is used. That is, the kernel design is penetration-proof—in particular to all those clever attacks that the kernel designers never contemplated.

This foundation of mathematical completeness raises the kernel design and evaluation process above a mere game of wits with an attacker; this is analogous to information theory as a foundation for modern cryptanalysis. A dramatic effect is that the kernel facilitates objective evaluation of internal security. The evaluator need not examine the nearly endless number of possible penetration attempts; he need only verify that the mathematical model is correctly implemented by the kernel. In other words, the kernel provides the verifiably reliable internal controls needed for multilevel security.

Figure 1. Secure computer system

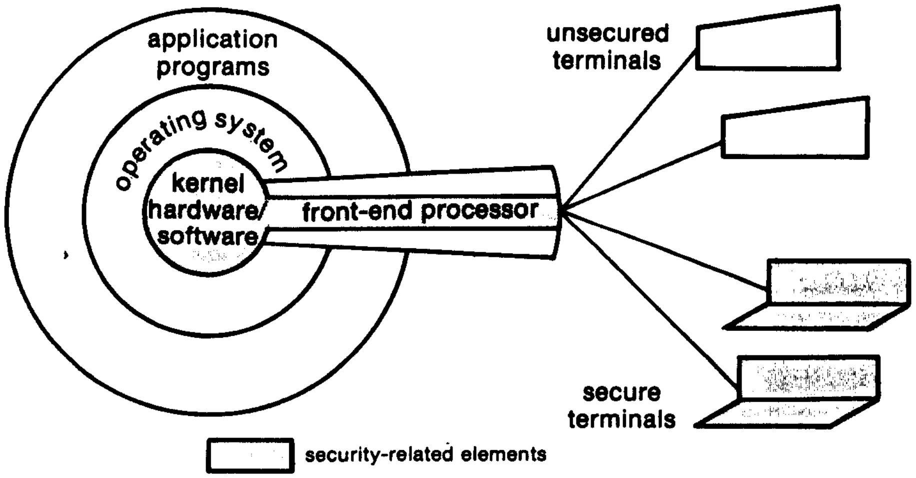

engineering feasibility

To be useful the kernel concept must be not only mathematically sound but also feasible to implement. Successful implementation is based on three engineering principles:

Completeness. A security kernel must be invoked on every access to data in the computer.

Isolation. A security kernel and its data base must be protected from unauthorized modification.

Verifiability. A security kernel must be sufficiently small and simple that its function can be completely tested and verified.

A laboratory security kernel for a commercial minicomputer (Digital Equipment Corporation model PDP-11/45) showed feasibility in 1974. The "virtual memory" hardware of this computer was a significant aid in ensuring the completeness and isolation of the kernel. This running kernel consisted of only about 1000 computer instructions. The experiment also established that it is much easier to introduce the kernel concept into an initial design than it is to retrofit it later.

The basis for the design (viz., kernel model) was mathematically verified. As with cryptographic devices, verification of the corresponding implementation was based more on careful engineering and extensive testing than on formal mathematics. Automated testing and program verification techniques indicated that the kernel implementation corresponded to the design. This laboratory prototype confirmed feasibility but was not oriented toward performance and efficiency evaluation. In passing, it is interesting to note that a tiger team tried and failed to penetrate its security.

performance

Performance was examined on a larger computer system. Negligible performance degradation (less than 1 percent) was experienced when the commercial Multics (for the Honeywell 6000 line) was modified to the kernel model. This Multics version was not implemented as a true kernel, i.e., the controls were distributed rather than collected into a small, verifiable entity; however, this version made all the security checks required in a kernel and thus confirmed that the kernel was not inherently inefficient.

The good security features of the kernel hardware were a major aid to performance, and these features are vendor-independent. The version was so successful that Honeywell included the resulting Access Isolation Mechanism in commercial Multics offerings for protection of privacy and business information. This system was used as the foundation for the terminated Air Force prototype; the prototype development was implementing a true, verifiable kernel.

functionality

A security kernel forces the computer user to be security-conscious but does not seriously degrade the capabilities of the computer. This was clearly demonstrated when the Multics modifications were successfully installed for those demanding users in the Pentagon: the constraints of the kernel design had minimal adverse impact on the users. Just as cryptography allows the secure use of standard commercial communication equipment, the kernel concept allows the secure use of standard commercial computer equipment and programs. The Pentagon facility with its classified processing confirmed the concepts for supporting a kernel-based computer in a total system security context.

Operational utility of the kernel was

further demonstrated with the initial minicomputer prototype. A demonstration showed the secure interface of operations and intelligence systems for fusion of tactical battlefield information. In addition, several military R&D efforts in various stages of completion have used major elements of the security kernel technology: a command and control network, a cryptographic controller, a nation-wide digital communication system, a large-scale "virtual machine monitor" system, a general-purpose minicomputer operating system, and a secure militarized minicomputer (based on the commercial Honeywell Level 6). Although they confirm the utility of the security kernel, none of these R&D efforts will lead to availability and operational use on a general basis.

security policy

Although the security kernel concept is not at odds with current policy, future policy must recognize and take advantage of kernel characteristics. Policy should recognize that the mathematical model provides a way to translate paper and pencil security rules into computer terms. In addition, a meaningful policy for multilevel mode would reflect the technological realities: either the entire system must be correct (not currently feasible) or else the security kernel must be used.

As with cryptographic devices, the kernel must be protected against subversion (e.g., insertion of a trap door) during its development. But protecting the kernel certainly involves far fewer people and a much more controlled environment than trying to protect all the computer programs of the system; thus, in contrast to contemporary systems, the kernel makes it tractable to protect against subversion. Furthermore, the evaluation (for certification) of internal computer security controls is a difficult technical task. The kernel approach to design and implementation makes such certification feasible, but this evaluation still requires highly capable

Table II. Commonality in security technology

	Cryptographic Mechanism	Security Kernel
threats negated rather than outlawed	wiretapping	penetration
standard commercial elements preserved	communications circuits	computers and programs
security sensitive portions limited	principally the crypto	principally the kernel
underlying basis precisely formulated	cryptographic algorithm	mathematical model
design evaluation criteria definitized	information theory	security theorems
implementation exactly meeting design	methodical engineering	verified programs
subversion controlled by physical security	manufacturing	programming
skilled experts needed for certification	cryptanalysts and engineers	computer scientists

technical experts—just as does the evaluation of cryptographic devices.

This approach conceptually parallels modern military cryptography. (See Table II.) Yet, development must be resumed and policy adjustments made if it is to be available on a general basis at any time in the immediate future. To be sure, there are competing demands for resources. Development of directly employable weapons (such as fighters) may always have higher priority than development of computer security, but as one observer put it: "How effective would those fighters be if plans for their employment were known in advance by an adversary who had penetrated the computer containing those plans?"[19] The security kernel is clearly the only currently available technology that can provide the security and operational capabilities we must have.

SECURITY often requires subjective judgments, and some may differ with the author on specific points. On balance it appears evident that a user who puts blind trust in the protection provided by computers for sensitive military information will seriously endanger security. In fact, most computers do not even include nominal features to support a military security system. Even when they do, the essence of the computer security problem is the technical efficacy of internal controls, and the evidence is clear that most internal controls are not dependable.

On the other hand, limiting computer use in order to avoid this problem is expensive and deprives us of vital operational capability. The effectiveness versus efficiency dilemma generates pressure for underestimating the threat and overconfidence in internal security controls. Unfortunately, these pressures have led the Air Force into a disturbing and increasing dependency on weak security controls even in the absence of evidence of effectiveness.

The Air Force recently terminated the single major DOD program for providing practical and scientifically sound internal controls—controls based on the security kernel concept. Past development has clearly demonstrated the feasibility, performance, and utility of this technology. However, because of lack of both a technical understanding and a meaningful policy, there is currently little official support for development of this promising capability.

Three basic actions must be taken to control the adverse impact of our computer security weakness:

- Promulgate a clear policy that distinguishes between dependence on external controls (dedicated mode) and internal controls (multilevel mode). It should not be possible to satisfy the policy without genuinely providing security. Multilevel mode without a technically sound basis should be expressly prohibited.
- Incorporate explicit military security controls in classified processing systems. These must be based on a precise specification of the required functions (as in the kernel model for the Pentagon Multics). This step is crucial to future introduction of multilevel security without complete system redesign. (In the interim this can also aid in the protection of privacy and valuable resources.)
- Resume security kernel development to provide technically sound multilevel security. As in the previous Air Force program, this should be oriented toward the competitive military acquisition process. Concurrently, policy must be changed to facilitate operational use of the kernel technology.

IT IS NOT easy to make a computer system secure, but neither is it impossible. The greatest error is to ignore the problem—a fatal mistake which obviously allows available solutions to remain unused. Failure in this one critical area introduces an Achilles' heel into our battlefield support systems—the cornerstone of the modern electronic Air Force.

Naval Postgraduate School
Monterey, California

Notes

1. Malcolm R. Currie, "Electronics: Key Military 'Force Multiplier,' " *Air Force Magazine*, July 1976, p. 44.
2. Edgar Ulsamer, "How ESD Is Building USAF's Electronic Eyes and Ears," *Air Force Magazine*, July 1977, p. 40.
3. Importance of electronics to the Air Force is indicated in "The Electronic Air Force," *Air Force Magazine*, July 1977, p. 29, which notes that this is the magazine's seventh annual issue devoted primarily to this "fundamental Air Force concern."
4. F. W. Winterbotham, *The Ultra Secret* (New York: Harper and Row, 1974), p. 11.
5. Ibid., p. 15.
6. Ibid., p. 107.
7. Ibid., p. 191.
8. David Kahn, *The Codebreakers* (New York: Macmillan Co., 1967), p. 67.
9. Winterbotham, p. 85.
10. Kahn, p. 591.
11. Ibid., p. 392.
12. Thomas Whiteside, "Dead Souls in the Computer," *The New Yorker*, August 29, 1977, pp. 59-62.
13. Tom Alexander, "Waiting for the Great Computer Rip-off," *Fortune*, July 1974, p. 143.
14. Bonnie Ginzburg, "Military Computers Easily Penetrable, AF Study Finds," *Washington Post*, August 8, 1976, p. A6.
15. In August 1976 Air Force Systems Command directed termination of the Electronic Systems Division's ADP System Security Program. Termination was completed by September 1977, halting development (that was proceeding well) of a secure general-purpose prototype to fully demonstrate operational acceptability and the associated development of specifications, policy recommendations, and evaluation criteria for general use.
16. Lawrence Curran, "Air Force 'Kernel' Attains Computer Security Using Existing Technology," *Electronics*, September 30, 1976, pp. 59, 61.
17. The author initially hypothesized the security kernel concept and its mathematical basis. Subsequent sponsored research at the MITRE Corporation completed the detailed formulation, as described in *ESD 1976 Computer Security Developments Summary*, MCI-76—2, Electronics Systems Division, Hanscom AFB, Massachusetts, January 1977.
18. W. L. Schiller, *The Design and Specification of a Security Kernel for the PDP-11/45*, ESD—TR-75-69, (Bedford, Massachusetts: MITRE Corporation, May 1975), p. 9.
19. "Computer Security: A Case of Priorities," *Electronics*, September 30, 1976, p. 10.

Technological progress has merely provided us with more efficient means for going backwards.

ALDOUS HUXLEY