

# An Overview of Computer Viruses in a Research Environment

*Matt Bishop*

Department of Mathematics and Computer Science  
Dartmouth College  
Hanover, NH 03755

## *ABSTRACT*

The threat of attack by computer viruses is in reality a very small part of a much more general threat, specifically attacks aimed at subverting computer security. This paper examines computer viruses as malicious logic in a research and development environment, relates them to various models of security and integrity, and examines current research techniques aimed at controlling the threats viruses in particular, and malicious logic in general, pose to computer systems. Finally, a brief examination of the vulnerabilities of research and development systems that malicious logic and computer viruses may exploit is undertaken.

## **1. Introduction**

A *computer virus* is a sequence of instructions that copies itself into other programs in such a way that executing the program also executes that sequence of instructions. Rarely has something seemingly so esoteric captured the imagination of so many people; magazines from *Business Week* to the *New England Journal of Medicine* [39][48][60][72][135], books [20][22][31][40][50][67][83][90][108][124], and newspaper articles [85][91][92][94][114][128] have discussed viruses, applying the name to various types of malicious programs.

As a result, the term “computer virus” is often misunderstood. Worse, many who do understand it do not understand protection in computer systems, for example believing that conventional security mechanisms can prevent virus infections, or are flawed because they cannot. But computer viruses use a number of well-known techniques in an unusual order; they do not employ previously-unknown methods. So, although existing computer security mechanisms were not designed specifically to counter computer viruses, many of those mechanisms were designed to deal with techniques used by computer viruses. While security mechanisms cannot prevent computer virus infections any more than they can prevent all attacks, they can impede a virus’ spread as well as make the introduction of a computer virus difficult, just as they can limit the damage done in an

---

This work was supported by grants NAG2-328 and NAG2-628 from the National Aeronautics and Space Administration to Dartmouth College.

attack, or make a successful attack very difficult. This paper tries to show the precise impact of many conventional security mechanisms on computer viruses by analyzing viruses in a general framework.

Because the probability of encountering a computer virus and the controls available to deal with it vary widely among different environments, this paper confines itself to that environment consisting of computers running operating systems designed for research and development, such as the UNIX<sup>1</sup> operating system, the VAX/VMS<sup>2</sup> operating system, and so forth. There is already a wealth of literature on computer viruses within the personal computing world (for example, see [34][62][65][124]), and a simple risk analysis (upon which we shall later elaborate) suggests that systems designed for accounting, inventory control, and other primarily business oriented operations are less likely to be attacked by using computer viruses than by other methods. So, while some of the following discussion may be fruitfully applied to computer systems in those environments (for example, see [1]), many of the underlying assumptions of system management and administration simply do not apply to those environments.

First, we shall review what a computer virus is, and analyze the properties that make it a threat to computer security. Next, we present a very brief history of computer viruses and consider whether their threat is relevant to research and development systems, and if so, how. After exploring some of the research in secure systems that show promise for coping with viruses, we examine several specific areas of vulnerability in research-oriented systems. We conclude with a quick summary.

## 2. What is a Computer Virus?

Computer viruses do not appear spontaneously [25]; an *attacker* must introduce one to the targeted computer system, usually by persuading, or tricking, someone with legitimate access into placing the virus on the system. This can readily be done using a *Trojan horse*, a program which performs a stated function while performing another, unstated and usually undesirable one (see sidebar 1).<sup>3</sup> For example, suppose a file used to boot a microcomputer contains a Trojan horse designed to erase a disk. When the microcomputer boots, it will execute the Trojan horse, which would erase the disk. Here, the overt function is to provide a basic operating system; the covert function is to erase the disk.

- 
1. UNIX is a registered trademark of AT&T Bell Laboratories.
  2. VAX and VMS are registered trademarks of Digital Equipment Corporation.
  3. D. Edwards first referred to this type of program as a "Trojan horse" in [4]

Many studies have shown the effectiveness of the Trojan horse attack (see [99][101], for example), and one such study [74] described a Trojan horse that reproduces itself (a *replicating Trojan horse*). If such a program *infects* another by inserting a copy of itself into the other file or process, it is a *computer virus*. (See sidebar 2; Leonard Adelman first called programs with the infection property “viruses” in a computer security seminar in 1983 [25].)

A computer virus infects other entities during its *infection phase*, and then performs some additional (possibly null) actions during its *execution phase*. Many view the infection phase as part of the “covert” action of a Trojan horse, and consequently consider the virus to be a form of the Trojan horse [44][69]. Others treat the infection phase as “overt” and distinguish between the virus and the Trojan horse, since a virus may infect and perform no covert action [25][97]. But all agree that a virus may perform covert actions during the execution phase.

Like Trojan horses [39], computer viruses are instances of *malicious logic* or *malicious programs*. Other programs which may be malicious but are not computer viruses are *worms*, which copy themselves from computer to computer<sup>4</sup>; *bacteria*, which replicate until all available resources of the host computer are absorbed; and *logic bombs*, which are run when specific conditions, such as the date being Friday the 13th, hold.

Malicious logic uses the user’s rights to perform their functions; a computer virus will spread only as the user’s rights will allow it, and can only take those actions that the user may take, since operating systems cannot distinguish between intentional and unintended actions. As the programs containing viruses are shared among users, the viruses spread among those users [25][97] until all programs writable by any infected program are themselves infected [56].

### **3. Malicious Logic, Computer Viruses, and Computer Security**

A site’s *security policy* describes how users may access the computer system or information on it, and the policy’s nature depends largely on how the system is to be used. *Military system security policies* deal primarily with disclosure of information, whereas *commercial security policies* deal primarily with the integrity of data on a system.

Security *mechanisms* that enforce policies partition the system into *protection domains* which define the set of objects that processes may access. *Mandatory access controls* prevent processes from crossing protection domain boundaries. *Discretionary access controls* condition permission to cross domain boundaries upon both the process identity and information associated with

---

4. Originally, a worm was simply a distributed computation [115]; it is now most often used in the above sense.

the object to be accessed.

Policies using mandatory access controls to prevent disclosure define a linear ordering of security levels, and a set of classes into which information is placed. Each entity's security classification is defined by the pair (*security level, set of classes*); the security classification of entity *A* *dominates* that of entity *B* if *A*'s security level is at least that of *B* and *A*'s set of classes contains all elements of *B*'s set of classes. Then the controls usually enforce some variant of the *Bell-LaPadula model* [9]: a subject may read an object only if the subject's security classification dominates that of the object (the *simple security property*) and a subject may modify an object only if the object's security classification dominates that of the subject (the *\*-property* or the *confinement property*). Hence subjects may obtain information only from entities with "lower" security classifications, and may disclose information only to entities with a "higher" security classification. These controls limit malicious logic designed to disclose information to the relevant protection domain; they do not limit malicious logic designed to corrupt information in "higher" security classifications.

Policies using discretionary access controls to limit disclosure assume that all processes of a given identity act with the authorization of that identity. When a program containing malicious logic is executed, the malicious logic executes with the same identity as that user's legitimate processes. The protection mechanism has no way to distinguish between acts done for the user and acts done for the attacker by the malicious logic.

Policies using mandatory access controls to limit modification of entities often implement the mathematical dual of the multilevel security model described above. Multilevel integrity models define integrity levels and classes analogous to those of the multilevel security models; then controls may enforce the *Biba integrity model* [11], which allows a subject to read an entity only if the entity's integrity classification dominates that of the subject (the *simple integrity property*), and a subject to modify an entity only if the subject's integrity classification dominates that of the entity (the *integrity confinement property*). This prevents a subject from modifying data or other programs at a higher integrity level, and a subject from relying on data or other programs at a lower integrity level. Hence, malicious logic can only damage those entities with lower or equal integrity classifications.

Lipner has proposed using the multilevel disclosure model to enforce multilevel integrity by assigning classifications and levels to appropriate user communities [87]; however, he notes that malicious logic could "write up" and thereby infect programs or alter production data and code. Clark and Wilson have proposed an alternate model [24] in which data and programs are manipu-

lated by well-defined “transformation procedures,” these procedures having been certified by the system security officer as complying with the site integrity policy. Hence computer viruses could only propagate among production programs if a transformation procedure which contains one is itself certified to conform to the integrity policy.

Policies using discretionary access controls to limit modification of entities make the same assumptions as security policies using discretionary access controls, with similar results.

Systems implementing multilevel security and integrity policies usually allow some small set of trusted entities to violate the stated policy when necessary for the smooth operation of the computer system. The usefulness of whatever security model the system implements depends to a very great extent on these exceptions; for should a trusted entity attempt to abuse its power to deviate from the strict policy, little can be done. The statements describing the effects of the controls on malicious logic above apply only to the model, and must be suitably modified for those situations in which a security policy allows (trusted) entities to violate the policy.

The two phases of a computer virus’ execution illustrate this. Infecting (altering) a program may be possible due to an allowed exception to the site’s integrity model. Executing a computer virus to disclose some information across protection domain boundaries may also be possible because of an allowed exception to the site’s disclosure model. So the virus may spread more widely because of the allowed exceptions.

An alternate view of malicious logic is that it causes the altered program to deviate from its specification. If this is considered an “error” as well as a breach of security, fault-tolerant computer systems, which are designed to continue reliable operation when errors occur, could constrain malicious logic. Designers of reliable systems place emphasis on both recovery and preventing failures [106]; however, if malicious logic discloses information or gives away rights, or controls other critical systems (such as life support systems), recovery may not be possible. So the areas of reliability and fault-tolerance are relevant to the study of malicious logic, but those areas of fault recovery are less so.

In the most general case, whether a given program will infect another is undecidable [2][25], so programs that look for virus infections must check characteristics of known viruses rather than rely on a general infection detection scheme. Further, viruses can be programmed to mutate, and hence be able to evade those agents, which in turn can be programmed to detect the mutations; and in the general case, whether or not one virus mutated to produce another virus is

also undecidable [30].

#### 4. A Brief History of Computer Viruses and Related Programs

One of the earliest documented replicating Trojan horses was a version of the game program *animal* which when played created another copy of itself. A later version deleted one copy of the first version, and then created *two* additional copies of itself. Because it spread even more rapidly than the first version, this later program supplanted the first entirely. After a preset date, whenever anyone played the second version, it deleted itself after the game ended [41].

Ken Thompson created a far more subtle replicating Trojan horse when he rigged a compiler to break login security [107][127]. When the compiler compiled the login program, it would secretly insert instructions to cause the resulting executable program to accept a fixed, secret password as well as a user's real password. Also, when compiling the compiler, the Trojan horse would insert commands to modify the login command into the resulting executable compiler. Thompson then compiled the compiler, deleted the new source, and reinstalled the old source. Since it showed no traces of being doctored, anyone examining the source would conclude the compiler was safe. Fortunately, Thompson took some pains to ensure that it did not spread further, and it was finally deleted when someone copied another version of the executable compiler over the sabotaged one. Thompson's point was that "no amount of source-level verification or scrutiny will protect you from using untrusted code" ([127], p. 763), which bears remembering, especially given the reliance of many security techniques relying on humans certifying programs to be free of malicious logic.

In 1983, Fred Cohen designed a computer virus to acquire privileges on a VAX-11/750 running UNIX; he obtained all system rights within half an hour on the average, the longest time being an hour, and the least being under 5 minutes. Because the virus did not degrade response time noticeably, most users never knew the system was under attack. In 1984 an experiment involving a UNIVAC 1108 showed that viruses could spread throughout that system too. Viruses were also written for other systems (TOPS-20<sup>5</sup>, VAX/VMS, and a VM/370<sup>6</sup> system) but testing their effectiveness was forbidden. Cohen's experiments indicated that the security mechanisms of those systems did little if anything to inhibit computer virus propagation [25][26].

In 1987, Tom Duff experimented on UNIX systems with a small virus that copied itself into executable files. The virus was not particularly virulent, but when Duff placed 48 infected pro-

---

5. TOPS-20 is a registered trademark of Digital Equipment Corporation.

6. VM/370 is a registered trademark of IBM.

grams on the most heavily used machine in the computing center, the virus spread to 46 different systems and infected 466 files, including at least one system program on each computer system, within eight days. Duff did not violate the security mechanisms in any way when he seeded the original 48 programs [45]. By writing another virus in a language used by a command interpreter common to most UNIX systems, he disproved a common fallacy [50] that computer viruses are intrinsically machine dependent, and cannot spread to systems of varying architectures.

On November 2, 1988, a program combining elements of a computer worm and a computer virus targeting Berkeley and Sun UNIX-based computers entered the Internet; within hours, it had rendered several thousand computers unusable [46][47][109][117][118][122][123][125]. Among other techniques, this program used a virus-like attack to spread: it inserted some instructions into a running process on the target machine and arranged for those instructions to be executed. To recover, these machines had to be disconnected from the network, rebooted, and several critical programs changed and recompiled to prevent re-infection. Worse, the only way to determine if the program had other malicious side effects (such as deleting files) was to disassemble it. Fortunately, its only purpose turned out to be to propagate. Infected sites were extremely lucky that the worm<sup>7</sup> did not infect a system program with a virus designed to delete files, or did not attempt to damage attacked systems. Since then, there have been several incidents involving worms [59][66][125].

In general, though, computer viruses and replicating Trojan horses have been laboratory experiments rather than attacks from malicious or careless users. This raises a question of risk analysis: do the benefits gained in defending against computer viruses offset the costs of recovery and the likelihood of being attacked?

As worded, the above question implies that the mechanisms defending against computer viruses are useful *only* against computer viruses. However, computer viruses use techniques that are also used in other methods of attack, such as scavenging<sup>8</sup>, as well as by other forms of malicious logic. Defenses which strengthen access controls to prevent illicit access, or which prevent or detect the alteration of other files, also limit, prevent, or detect these other attacks as well. So, a more appropriate question is whether the benefits gained in defending against all such attacks offset the costs of recovery and the likelihood of being attacked.

---

7. We use the conventional terminology of calling this program a “computer worm” because its dominant method of propagation was from computer system to computer system. Others, notably [46], have labelled it a “computer virus” using a taxonomy more firmly grounded in biology than the conventional one.

8. Reading private files to obtain information (such as user names and passwords) that can then be used to break into other systems, or other parts of the system on which the information is found.

Because this paper focuses primarily on computer viruses, we shall not delve into the history of computer security or malicious logic in general. Suffice it to say that the vulnerability of computer systems to such attacks is well known, and attacks on computer systems are common enough (see both [99] and [101] for descriptions of such incidents) that the use of mechanisms to inhibit them is generally agreed to be worthwhile.

## **5. Current Research in Malicious Logic and Computer Viruses**

The effectiveness of any security mechanism depends upon the security of the underlying base on which the mechanism is implemented, and the correctness of the necessary checking done at each step. If the trust in the base or in the checking is misplaced the mechanism will not be secure. Thus “secure” is a relative notion, as is “trust,” and mechanisms to enhance computer security attempt to balance the cost of the mechanism with the level of security desired and the degree of trust in the base that the site accepts as reasonable. Research dealing with malicious logic assumes the interface, software, and/or hardware used to implement the proposed scheme performs exactly as desired, meaning the trust is in the underlying computing base, the implementation, and (if done) the verification.

Current research uses specific properties of computer viruses to detect and limit their effects. Because of the fundamental nature of these properties, these defenses work equally well against most other forms of malicious logic.

### **5.1. Computer Viruses Acting as Both Data and Instructions**

Techniques exploiting this property treat all programs as type “data” until some certifying authority changes the type to “executable” (instructions). Both new systems designed to meet strong security policies and enhancements to existing systems use this method.

Boebert and Kain [18] have proposed labelling subjects and objects in the Logical Coprocessor Kernel or LOCK (formerly the Secure Ada Target or SAT) [17][61][112][113], a system designed to meet the highest level of security under the Department of Defense criteria [43]. Once compiled, programs have the label “data,” and cannot be executed until a sequence of specific, auditable events changes the label to “executable.” After that, the program cannot be modified. This scheme recognizes that viruses treat programs as data (when they infect them by changing the file’s contents) and as instructions (when the program executes and spreads the virus), and rigidly separates the two. The Argus Security Model [3] uses the same principle.

Duff [45] has suggested a variant for UNIX-based systems. Noting that users with execute permission for a file usually also have read permission, he proposes that files with execute permission be of type “executable,” and those without it be of type “data.” Unlike the LOCK, “executable” files could be modified but doing so would change the type to “data.” If the certifying authority were the omnipotent user, the virus could spread only if run as that user. To prevent infection from non-executable files, libraries and other system components of programs must also be certified before use.

Both the LOCK scheme and Duff’s proposal trust that the administrators will never certify a program containing malicious logic (either by accident or deliberately), and that the tools used in the certification process are not themselves corrupt.

## **5.2. Viruses Assuming the Identity of a User**

Among the many enhancements to discretionary access controls are suggestions to allow the user to reduce the associated protection domain [29][72][121][134]; to base access to files on some characteristic of the command or program [27][81], possibly including subject authorizations as well [25]; and to use a knowledge-based subsystem to determine if a program makes reasonable file accesses [73]. Allowing users to specify semantics for file accesses [10][36] may prove useful in some contexts, for example protecting a limited set of files.

All such mechanisms trust the users to take explicit action to limit their protection domains sufficiently; or trust tables to describe the programs’ expected actions sufficiently for the mechanism to apply those descriptions, and the mechanism to handle commands with no corresponding table entries effectively; or they trust specific programs and the kernel, when those would be the first programs a virus would attack.

## **5.3. Viruses Crossing Protection Domain Boundaries by Sharing.**

Inhibiting users in different protection domains from sharing programs or data will inhibit viruses from spreading among those domains. For example, when users share procedures, the LOCK keeps only one copy of the procedure in memory. A master directory, accessible only to a trusted hardware controller, associates with each procedure a unique owner, and with each user a list of others whom that user trusts. Before executing any procedure, the dynamic linker checks that the user executing the procedure trusts the procedure’s owner [16]. This scheme assumes that users’ trust in one another is always well-placed.

A more general proposal [137] suggests placing programs to be protected at the lowest possible level of an implementation of a multilevel security policy. Since the mandatory access controls will prevent those processes from writing to objects at lower levels, any process can read the programs but no process can write to them. Such a scheme would have to be combined with an integrity model to provide protection against viruses to prevent both disclosure and file corruption. Carrying this idea to its extreme would result in isolation of each domain; since sharing is not possible, no viruses can propagate. Unfortunately, the usefulness of such systems would be minimal.

#### 5.4. Viruses Altering Files

Mechanisms using *manipulation detection codes* (or *MDCs*) apply some function to a file to obtain a set of bits called the *signature block* and then encrypt that block. If, after recomputing the signature block and reencrypting it, the result differs from the stored signature block, the file has changed [86][95], possibly due to infection or some other cause not related to viruses.

An assumption is that the signed file does not contain a virus before it is signed. Page [100] has suggested expanding the model in [17] to include the software development process (in effect limiting execution domains for each development tool and user) to ensure software is not contaminated during development. Pozzo and Grey [104][105] have implemented Biba's integrity model on the distributed operating system LOCUS [103] to make the level of trust in the above assumption explicit. They have different classes of signed executable programs. *Credibility ratings* (Biba's "integrity levels") assign a measure of trustworthiness on a scale of 0 (unsigned) to  $N$  (signed and formally verified), based on the origin of the software. Trusted file systems contain only signed executable files with the same credibility level. Associated with each user (subject) is a *risk level* that starts out as the highest credibility level. Users may execute programs with credibility levels no less than their risk level; when the credibility level is lower than the risk level, a special "run-untrusted" command must be used.

All integrity-based schemes rely on software which if infected may fail to report tampering. Performance will be affected as encrypting the file or computing the signature block may take a significant amount of time. The encrypting key must also be secret, for if not then malicious logic can easily alter a signed file without the change being detected.

Network implementations of MDC-based mechanisms require that public keys be certified by a trusted authority and distributed in a trusted fashion (see for example [15][75]). If the key distribution mechanism used the same paths as the data transmission and the public keys were not ver-

ifiable using an out-of-bands method, a malicious site (or set of cooperating malicious sites) could alter the data or program being sent, recompute the signature block and sign it with its own (bogus) private key, and then transmit the data; when the public key were requested, it would simply send the one corresponding to the (bogus) private key. The more general (non-network) software distribution problem has similar requirements [35].

Anti-virus agents check files for specific viruses and if present either warn the user or attempt to “cure” the infection by removing the virus. Many such agents exist for personal computers, but since each must look for a particular virus or set of viruses, they are very specific tools and, because of the undecidability results stated earlier, cannot deal with viruses not yet analyzed.

### **5.5. Viruses Performing Actions Beyond Specification**

Fault-tolerant techniques keep systems functioning correctly when the software or hardware fails to perform to specification. Joseph and Avizienis have suggested treating a virus’ infection and execution phases as errors. The first such proposal [70][71] breaks programs into sequences of non-branching instructions, and checksums each sequence, storing the results in encrypted form. When the program is run, the processor recomputes checksums, and at each branch, a co-processor compares the computed checksum to the encrypted checksum; if they differ, an error (which may be an infection) has occurred. Later proposals advocate checking each instruction [35]. These schemes raise issues of key management and protection, as well as how much the software managing keys, transmitting the control flow graph to the co-processor, and implementing the recovery mechanism, may be trusted.

A proposal based on *N-Version Programming* [5] requires implementing several different versions of an algorithm, running them concurrently and periodically checking intermediate results against each other. If they disagree, the value assumed correct is the intermediate value that a majority of the programs have obtained, and the programs with a different value are malfunctioning (possibly due to malicious logic). This requires a majority of the programs not to be infected, and the underlying operating system to be secure. Also, the issue of the efficacy of N-version programming is highly questionable [77]. Despite claims that the method is feasible [6][23], detecting the spread of a virus would require voting upon each file system access; to achieve this level of comparison, the programs would all have to implement the same algorithm, which defeats the purpose of using N-version programming [78].

## 5.6. Viruses Altering Statistical Characteristics

Proposals to examine the appearance of programs for identical sequences of instructions or byte patterns [69][137] require a high number of comparisons and would need to take into account the reuse of common library routines or of code [76]. Malicious logic might be present if a program appears to have more programmers than were known to have worked on it, or if one particular programmer appears to have worked on many different and unrelated programs [137]; but several assumptions must first be validated, namely that programmers have their own individual styles of writing programs, that the executable programs generated by the compilers will reflect these styles, and that a *coding style analyzer* can distinguish these styles from one another. If an object file contains conditionals not corresponding to any in the source, the object may be infected [54]. A fourth proposal suggests designing a filter to detect, analyze, and classify all modifications that a program will make as ordinary or suspicious [32].

Finally, Dorothy Denning has suggested using an intrusion-detection expert system to detect viruses by looking for increases in the size of files, increases in the frequency of writing to executable files, or alterations in the frequency of executing a specific program in ways not matching the profile of users spreading the infection [38]. Several such systems have been implemented [8][88][126] and have detected many anomalies without noticeably degrading the monitored computer. These experiments did not attempt to validate claims about detecting viruses.

Those research proposals that are being implemented are either targeted for specific architectures or are in the very early stages of development. This state of affairs is unsettling for the managers and administrators of existing systems, who need to take some action to protect their users and systems.

## 6. Vulnerabilities of Existing Research-Oriented Systems

The vulnerabilities exploited by a computer virus can also be exploited by other forms of malicious logic, and unless the purpose of the attack is to cause mischief, the other forms of malicious logic are much easier to create. Rather than describe appropriate countermeasures, we simply note that these will differ from environment to environment, and no such list (or even set of lists) can accurately reflect the idiosyncracies of all the different research and development systems and environments; in short, providing such a generic list could give a very false sense of security.

This section discusses the areas of vulnerability. While we emphasize computer viruses throughout, these same vulnerabilities can be exploited by Trojan horses, computer worms, other

forms of malicious logic, and, more generally, other types of attacks. We leave it to the reader to formulate appropriate techniques to detect or hinder attacks exploiting each area. (Sidebar 3 offers a starting point for UNIX-based systems.)

## **6.1. Computing Base**

Users assume that the computer system provides a set of trustworthy tools for compiling, linking and loading, and running programs. In most systems, the “trust” is the user’s estimate of the quality of the tools available [28] and the working environment. If the estimates are incorrect, the system may be subverted.

Even systems with security enhancements are vulnerable. One version of the UNIX operating system with security enhancements was breached when a user created a version of the directory lister, with a Trojan horse, in his home directory. He then requested assistance from the system operator, who changed to the user’s home directory, and listed the names of the files in it. As the command interpreter checked for commands in the current working directory and *then* in the system directories, the user’s doctored lister, not the system lister, was executed [120].

In the above, the system administrator trusted the command interpreter to look for system programs before executing programs in users’ directories. Other examples include trusting that the login banner being presented is actually from the login program and not from a user’s program which will record passwords [58], or that page faults cannot be detected while checking passwords one character at a time [82].

## **6.2. Sharing Hardware and Software**

Intimately bound with the notion of trust is the ability to share. When many computers share a copy of an infected program, every file accessible from every one of those machines can be infected. Methods of sharing include making and distributing copies of software, accessing bulletin board systems, public file servers, and obtaining source files from remote hosts using a network or electronic mail.

The probability of any new program containing malicious logic depends on the integrity of the author (or authors), the security and integrity of the computer on which they worked, on which the distribution was prepared, and on the method of distribution. Programs sent through electronic mail or posted to bulletin boards may be altered in transit, either by someone modifying them while they sit on an intermediate node, or while they are crossing networks [133]. Further, electronic

messages can easily be forged [116][132], so it is unwise to rely on such a program's stated origin.

In the early 1980s a program posted to the USENET news network contained a command to delete all files on the system in which it was run. Some system administrators executed the program with unlimited privileges, thereby damaging their systems. In another case, although vendors usually take care that their software contains no malicious logic, a company selling software for the Macintosh<sup>9</sup> unwittingly delivered copies of programs infected by a computer virus which printed a message asking for universal peace [51].

### 6.3. Integrity of Programs

The infection phase of a virus' actions require writing to files; for reasons discussed earlier, discretionary access controls provide little protection. Typically some form of auditing is used to detect changes [14][19]; however, auditing schemes cannot prevent damage, but only attempt to provide a record of it and (possibly) indicate the culprit. The best auditing methods use a mechanism that records changes to files or their characteristics. Such schemes require kernel modifications [102] and should be designed into new systems [57][79][96]; if a site has only object code, it cannot add these mechanisms and so must scan the file system [13]. Audit logs must also be protected from illicit modification; again, an element of trust in the underlying subsystem is needed.

A computer virus can defeat any auditing scheme by infecting a file and then altering the file's contents or characteristics during the audit, for example by restoring the uncorrupted version temporarily. An example of such a *stealth virus* is the 4096 (personal computer) virus [89].

No program can determine if an arbitrary virus has infected a file because of the undecidability results cited earlier; however, *virus detectors* or *anti-virus agents* can check files for specific virus. If a virus detector reports that no infection is present, the file may contain a virus unknown to the detector, or the detector may be corrupt. In February 1989, at Dartmouth College, a user ran an infected version of the virus detection program Interferon, infecting files on his disk. More widely known is the Trojan horse in a doctored copy of the anti-virus program FLUSHOT [64]; later versions are called FSP+ to avoid confusion with the tampered version [7].

### 6.4. Backups and Recovery

Using backups to replace infected files, or files which contain malicious logic, may remove such programs from the system. As most systems make backup copies of files which have changed

---

9. Macintosh is a Registered Trademark of Apple Computer

since the time the previous backup was made, it is quite likely that several backups will need to be examined to find an uncontaminated version of the infected program. Further, unless all malicious programs are found and restored at the same time, the restoration of some uncorrupted programs may do little (for example, computer viruses still resident on the system could infect the newly-restored programs).

If the backup and restore programs themselves contain malicious logic that prevents uncorrupted software from being restored, then the backups are useless until a way is found to replace (or fix) the restore program. Worse, some research and development systems (such as variants of the UNIX operating system) do not allow users to “lock” devices, so one user can access media mounted by another user. Thus, between the mounting and the attempt to restore, another program containing malicious logic could easily infect or erase a mounted backup.

## **6.5. The Human Factor**

It has been said that computer viruses are a management issue, because they are introduced by people [37]; the same may be said for all malicious logic, and computer security in general. Ideally, security procedures should balance the security and safety of the system and data with the needs of the users and systems personnel to get work done. All too often, users (and systems personnel) see them as burdens to be evaded. Lack of awareness of the reasons for security procedures and mechanisms leads to carelessness or negligence, which can in turn lead to system compromise (see for example [101]).

Little if anything can be done to prevent compromise by trusted personnel. Malicious users and system administrators can often circumvent security policy restrictions without being stopped, or even detected, by using the exceptions to the mechanisms enforcing the policies. (See [99] for examples of these “inside jobs.”) The study of computing ethics, or of a code of ethical conduct, reduces this threat by making clear what actions are considered acceptable; should a breach occur, legal remedies may be available [55][111].

## **6.6. Multiple Levels of Privilege**

Multi-user computer systems often provide many different levels of privilege; for example, UNIX provides a separate set of privileges for each user, and one all-powerful *superuser*. Enforcing the *principle of least privilege* [110] can limit the files that malicious logic can read or write.

If someone using a privileged account accidentally executes a program containing a com-

puter virus, the virus will spread throughout the system rapidly [45]. Hence, simply logging in as a privileged user and remaining so empowered increases the possibility of accidentally triggering some form of malicious logic. More subtle is the use of programs which can cross protection domain boundaries; when the boundary being crossed involves the addition of a privilege or capability that enables the user to affect objects in many other protection domains (such as changing from an unprivileged to a privileged mode), a malicious program could read or alter data or programs not normally accessible to the user. In general, computer systems do not force such programs to function with as few privileges as possible. For example, the `setuid` and `setgid` mechanism of UNIX [12][21][84] violate this principle.

A related but widely-ignored problem is the use of “smart” terminals to access privileged accounts. These terminals will respond to control sequences from a host by transmitting portions of the text on their screen back to the host [52], and often perform simple editing functions for the host. Such a terminal can issue a computer virus’ commands in the name of the terminal’s user when appropriate text and control sequences are sent to it (for example, by using an inter-terminal communications program or displaying files with appropriate characters in it.) These commands could instruct the computer to execute an infected program, which would run in the protection domain of the user of the terminal (and not that of the attacker). As many computers use such terminals as their consoles, and allow access to the most privileged accounts only when the user is at the console, the danger is obvious.

### **6.7. Direct Device Access**

The *principle of complete mediation* [110] requires checking the validity of every access. Although multi-user systems have virtual memory protection to prevent processes from writing into each other’s memory, some represent devices and memory as addressable objects (such as files). If these objects are improperly or inadequately protected, a process could bypass the virtual memory controls and write to any location in memory by placing data and addresses on the bus, thereby altering the instructions and data in another’s memory space (the “core war” games [42] did this). If any process could write to disks without the kernel’s intervention, anyone can change executable programs regardless of their protection – and a virus can easily spread by taking advantage of the (lack of) protection.

## **7. Conclusion**

This paper has described the threats that computer viruses pose to research and develop-

ment multi-user computer systems; it has attempted to tie those programs with other, usually simpler, programs that can have equally devastating effects. Although reports of malicious programs in general abound, no non-experimental computer viruses have been reported on mainframe systems.<sup>10</sup> Noting that the number of people with access to mainframes is relatively small compared to the number with access to personal computers [130], Highland suggests that as malicious people make up a very small fraction of all computer programmers, most likely fewer malicious people use research and development systems than personal computers [64]. A more persuasive argument, advanced by Fåk [49] and supported by Kurzban [80] is that, as only programmers can create computer viruses, and malicious mainframe programmers can accomplish their goals with less trouble than writing a computer virus, computer virus attacks will most likely be confined to personal computers. Exceptions would most likely be motivated by a perceived intellectual challenge of creating a virus, by a desire to demonstrate limits of existing security mechanisms, by a desire for publicity, or attacks launched simply by carelessness or error [98].<sup>11</sup>

Should an attacker use a computer virus or other malicious program, security mechanisms currently in use will be as effective as they are against other types of attacks. As with attempts to breach security in general, though, people can prepare for such an attack and minimize the damage done. This paper has described several vulnerabilities in the research and development environment that malicious programs could exploit, and also discussed research underway to improve defenses against malicious logic. How effective these new mechanisms will be in reducing the vulnerabilities, only time will tell.

**Acknowledgments:** Thanks to Holly Bishop, Ken Bogart, André Bondi, Emily Bryant, Peter Denning, Donald Johnson, John Rushby, Eugene Spafford, Ken Van Wyk, and the anonymous referees, all of whose comments and advice improved the quality of the paper greatly. Josh Alden of the Dartmouth Virus Clinic described the Interferon infection incident, Robert Van Cleef and Gene Spafford helped reconstruct the USENET logic bomb incident, and Ken Thompson confirmed that he had indeed doctored an *internal* version of the C compiler as described in [127]. My thanks to them also.

---

10. Cohen tantalizingly claims that one has been found, but reports no other details [27]. Suppression of details (or, more commonly, the existence) of attacks, virus or otherwise, is common; it is estimated that victims report only 10% to 35% of computer crimes in general [119][129], in part to prevent embarrassment or loss of public confidence in the company, or to avoid the expense of gathering sufficient evidence to prosecute the offender [101].

11. It is worth noting that the author of the Internet worm stated that the worm disabled machines due to a programming error [93].

## References

- [1] G. Al-Dossary, "Computer Virus Prevention and Containment on Mainframes," *Computers and Security* **9**(2) (Apr. 1990) pp. 131-137.
- [2] L. Adelman, "An Abstract Theory of Computer Viruses," *Advances in Cryptology – CRYPTO '88 Proceedings*, Springer-Verlag, New York, NY (Aug. 1988) pp. 354-374.
- [3] M. Adkins, G. Dolsen, J. Heaney, and J. Page, "The Argus Security Model," *Twelfth National Computer Security Conference Proceedings* (Oct. 1989) pp. 123-134.
- [4] J. Anderson, "Computer Security Technology Planning Study," ESD-TR-73-51, Air Force Electronic Systems Division, Hanscom Air Force Base, MA (1974).
- [5] A. Avizienis, "The N-Version Approach to Fault-Tolerant Software," *IEEE Transactions on Software Engineering* **SE-11**(12) (Dec. 1985) pp. 1491-1501.
- [6] A. Avizienis, M. Lyu, and W. Schutz, "In Search of Effective Diversity: A Six-Language Study of Fault-Tolerant Control Software," Technical Report CSD-870060, University of California, Los Angeles, CA (Nov. 1987).
- [7] D. Bader, "Bad Versions of FLUSHOT (for IBM PC)," *Virus-L Digest* **1**(8) (Nov. 15, 1988).
- [8] D. Bauer and M. Koblenz, "NDIX – A Real-Time Intrusion Detection Expert System," *1989 Summer USENIX Conference Proceedings* (June 1988) pp. 261-274.
- [9] D. Bell and L. LaPadula, "Secure Computer Systems: Unified Exposition and MULTICS Interpretation," Technical Report MTR-2997, MITRE Corporation, Bedford, MA (July 1975).
- [10] B. Bershad and C. Pinkerton, "Watchdogs: Extending the UNIX File System," *1988 Winter USENIX Conference Proceedings* (Feb. 1988) pp. 267-276.
- [11] K. Biba, "Integrity Considerations for Secure Computer Systems," Technical Report ESD-TR-76-372, Air Force Electronic Systems Division, Hanscom Air Force Base, MA (1977).
- [12] M. Bishop, "How to Write a Setuid Program," *login*: **12**(1) (Jan. 1987) pp. 5-11.
- [13] M. Bishop, "Auditing Files on a Network of UNIX Machines," *Proceedings of the UNIX Security Workshop* (Aug. 1988) pp. 51-52.
- [14] M. Bishop, "A Model of Security Monitoring," *Proceedings of the Fifth Annual Computer Security Applications Conference* (Dec. 1989) pp. 46-52.

- [15] M. Bishop, "An Authentication Mechanism for USENET," *1991 Winter USENIX Conference Proceedings* (Jan. 1991) pp. 281-287.
- [16] W. Boebert and C. Ferguson, "A Partial Solution to the Discretionary Trojan Horse Problem," *Proceedings of the Eighth Computer Security Conference* (sep. 1985) pp. 245-253.
- [17] W. Boebert and R. Kain, "A Practical Alternative to Hierarchical Integrity Policies," *Proceedings of the Eighth Computer Security Conference* (Sep. 1985) pp. 18-27.
- [18] W. Boebert, W. Young, R. Kain, and S. Hansohn, "Secure Ada Target: Issues, System Design, and Verification," *Proceedings of the 1985 Symposium on Security and Privacy* (Apr. 1985) pp. 176-183.
- [19] D. Bonyun, "The Role of a Well Defined Auditing Process in the Enforcement of Privacy Policy and Data Security," *Proceedings of the 1981 Symposium on Security and Privacy* (Apr. 1981) pp. 19-25.
- [20] J. Brunner, *The Shockwave Rider*, Ballantine York City, NY (1975).
- [21] S. Bunch, "The Setuid Feature in UNIX and Security," *Tenth National Computer Security Conference Proceedings* (Sep. 1987) pp. 245-253.
- [22] R. Burger, *Computer Viruses – A High-Tech Disease*, Abacus, Grand Rapids, MI (1988).
- [23] L. Chen, "Improving Software Reliability by N-Version Programming," Technical Report Eng-7843, University of California, Los Angeles, CA (Aug. 1978).
- [24] D. Clark and D. Wilson, "A Comparison of Commercial and Military Computer Security Policies," *Proceedings of the 1987 Symposium on Security and Privacy* (Apr. 1987) pp. 184-194.
- [25] F. Cohen, "Computer Viruses: Theory and Experiments," *Seventh DOD/NBS Computer Security Conference Proceedings* (Sep. 1984) pp. 240-263.
- [26] F. Cohen, "Computer Viruses: Theory and Experiments," *Computers and Security* **6**(1) (Feb. 1987) pp. 22-35.
- [27] F. Cohen, "On the Implications of Computer Viruses and Methods of Defense," *Computers and Security* **7**(2) (Apr. 1988) pp. 167-184.
- [28] F. Cohen, "Maintaining a Poor Person's Information Integrity," *Computers and Security* **7**(5) (Oct. 1988) pp. 489-494.
- [29] F. Cohen, "Practical Defenses Against Computer Viruses," *Computers and Security* **8**(2)

- (Apr. 1989) pp. 149-160.
- [30] F. Cohen, "Computational Aspects of Computer Viruses," *Computers and Security* **8**(4) (June 1989) pp. 325-344.
- [31] F. Cohen, *A Short Course on Computer Viruses*, ASP Press, Pittsburgh, PA (1990).
- [32] S. Crocker and M. Pozzo, "A Proposal for a Verification-Based Virus Filter," *Proceedings of the 1989 IEEE Symposium on Security and Privacy* (May 1989) pp. 319-324.
- [33] D. Curry, "Improving the Security of Your UNIX System," Technical Report ITSTD-721-FR-90-91, SRI International, Menlo Park, CA 94025 (Apr. 1990).
- [34] J. David, "Treating Viral Fever" *Computers and Security* **7**(2) (Apr. 1988) pp. 255-258.
- [35] G. Davida, Y. Desmedt, and B. Matt, "Defending Systems Against Viruses through Cryptographic Authentication," *Proceedings of the 1989 Symposium on Security and Privacy* (May 1989) pp. 312-318.
- [36] G. Davida and B. Matt, "UNIX Guardians: Delegating Security to the User," *Proceedings of the UNIX Security Workshop* (Aug. 1988) pp. 14-23.
- [37] H. DeMaio, "Viruses – Management Issue," *Computers and Security* **8**(5) (Oct. 1989) pp. 381-388.
- [38] D. Denning, "An Intrusion-Detection Model," *IEEE Transactions on Software Engineering* **SE-13**(2) (Feb. 1987) pp. 222-232.
- [39] P. Denning, "The Science of Computing: Computer Viruses," *American Scientist* **76**(3) (May 1988) pp. 236-238.
- [40] P. Denning, *Computers Under Attack: Intruders, Worms, and Viruses*, Addison-Wesley Publishing Co., Reading, MA (1990),
- [41] A. Dewdney, "Computer Recreations: A Core War Bestiary of Viruses, Worms, and Other Threats to Computer Memories," *Scientific American* **252**(3) (Mar. 1985) pp. 14-23.
- [42] A. Dewdney, "Computer Recreations," *Scientific American* **256**(1) (Jan. 1987) pp. 14-20.
- [43] *Trusted Computer System Evaluation Criteria*, DOD 5200.28-STD, Department of Defense (Dec. 1985).
- [44] D. Downs, J. Rub, K. Kung, and C. Jordan, "Issues in Discretionary Access Control," *Proceedings of the 1984 IEEE Symposium on Security and Privacy* (Apr. 1984) pp. 208-218.

- [45] T. Duff, "Experiences with Viruses on UNIX Systems," *Computing Systems* **2**(2) (Spring 1989) pp. 155-172.
- [46] M. Eichin and J. Rochlis, "With Microscope and Tweezers: An Analysis of the Internet Virus of November 1988," *Proceedings of the 1989 IEEE Symposium on Security and Privacy* (Apr. 1989) pp. 326-343.
- [47] T. Eisenberg, D. Gries, J. Hartmanis, D. Holcomb, M. Lynn, and T. Santoro, *The Computer Worm: A Report to the Provost of Cornell University on an Investigation Conducted by the Commission of Preliminary Enquiry*, Cornell University, Ithaca, NY (Feb. 1989).
- [48] P. Elmer-DeWitt, "Invasion of the Data Snatchers: A Virus Epidemic Strikes Terror in the Computer World," *Time* (Sep. 26, 1988) pp. 62-67.
- [49] V. Fåk, "Are We Vulnerable to a Virus Attack: A Report from Sweden," *Computers and Security* **7**(2) (Apr. 1988) pp. 151-155.
- [50] R. Farrow, *UNIX System Security*, Addison-Wesley Publishing Co., Reading, MA (1991).
- [51] P. Fites, P. Johnston, and M. Kratz, *The Computer Virus Crisis*, Van Nostrand Reinhold, New York City, NY (1988).
- [52] M. Gabriele, "'Smart' Terminals for Trusted Computer Systems," *Ninth National Computer Security Conference Proceedings* (Sep. 1986) pp. 16-20.
- [53] S. Garfinkel and G. Spafford, *Practical UNIX Security*, O'Reilly and Associates (1991).
- [54] P. Garnett, "Selective Disassembly: A First Step Towards Developing a Virus Filter," *Fourth Aerospace Computer Security Conference* (Dec. 1988) pp. 2-6.
- [55] M. Gemignani, "Viruses and Criminal Law," *CACM* **32**(6) (June 1989) pp. 669-671.
- [56] W. Gleissner, "A Mathematical Theory for the Spread of Computer Viruses," *Computers and Security* **8**(1) (Feb. 1989) pp. 35-41.
- [57] V. Gligor, C. Chandrasekaran, R. Chapman, L. Dotterer, M. Hecht, W. Jiang, A. Johri, G. Luckenbaugh, and N. Vasudevan, "Design and Implementation of Secure Xenix," *IEEE Transactions on Software Engineering* **SE-13**(2) (Feb. 1987) pp. 208-220.
- [58] F. Grampp and R. Morris, "UNIX Operating System Security," *AT&T Bell Laboratories Technical Journal* **63**(8) (Oct. 1984) pp. 1649-1672.
- [59] J. Green and P. Sisson, "The 'Father Christmas' Worm," *Twelfth National Computer Security Conference Proceedings* (Oct. 1989) pp. 359-368.

- [60] K. Hafner, "Is Your Computer Secure?," *Business Week* (Aug. 1, 1987) pp. 64-72.
- [61] J. Haigh and W. Young, "Extending the Non-Interference Version of MLS for SAT," *Proceedings of the 1986 IEEE Symposium on Security and Privacy* (Apr. 1986) pp. 232-239.
- [62] H. Highland, "Random Bits and Bytes: Case History of a Virus Attack," *Computers and Security* **7**(1) (Feb. 1988) pp. 3-5.
- [63] H. Highland, "Random Bits and Bytes: Case History of a Virus Attack," *Computers and Security* **7**(1) (Feb. 1988) pp. 6-7.
- [64] H. Highland, "Random Bits and Bytes: Computer Viruses – A Post-Mortem," *Computers and Security* **7**(2) (Apr. 1988) pp. 117-127.
- [65] H., Highland, "The Brain Virus: Fact and Fantasy," *Computers and Security* **7**(4) (Aug. 1988) pp. 367-370.
- [66] H. Highland, "Random Bits and Bytes: Another Poor Password Disaster," *Computers and Security* **9**(1) (Feb. 1990) p. 10.
- [67] L. Hoffman, *Rogue Programs: Viruses, Worms, and Trojan Horses*, Van Nostrand Reinhold, New York City, NY (1990).
- [68] Homer, *The Odyssey*, Penguin Books, New York City, NY (1946).
- [69] H. Israel, "Computer Viruses: Myth or Reality?," *Tenth National Computer Security Conference Proceedings* (Sep. 1987) pp. 226-230.
- [70] M. Joseph, "Towards the Elimination of the Effects of Malicious Logic: Fault Tolerance Approaches," *Tenth National Computer Security Conference Proceedings* (Sep. 1987) pp. 238-244.
- [71] M. Joseph and A. Avižienis, "A Fault Tolerant Approach to Computer Viruses," *Proceedings of the 1988 Symposium on Security and Privacy* (Apr. 1988) pp. 52-58.
- [72] J. Juni and R. Ponto, "Computer-Virus Infection of a Medical Diagnostic Computer," *New England Journal of Medicine* **320**(12) (Mar. 12, 1989) pp. 811-812.
- [73] P. Karger, "Limiting the Damage Potential of Discretionary Trojan Horses," *Proceedings of the 1987 Symposium on Security and Privacy* (Apr. 1987) pp. 32-37.
- [74] P. Karger and R. Schell, "MULTICS Security Evaluation: Vulnerability Analysis," Technical Report ESD-TR-74-193, Air Force Electronic Systems Division, Hanscom Air Force Base, MA (1974).

- [75] S. Kent and J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part II -- Certificate-Based Key Management*, RFC 1114 (Aug. 1989).
- [76] B. Kernighan and T. Plauger, *The Elements of Programming Style*, McGraw-Hill Book Co., New York City, NY (1974).
- [77] J. Knight and N. Leveson, "An Experimental Evaluation of the Assumption of Independence in Multi-version Programming," *IEEE Transactions on Software Engineering* **SE-12**(1) (Jan. 1986) pp. 96-109.
- [78] J. Knight and N. Leveson, "On N-version Programming," *Software Engineering Notes* **15**(1) (Jan. 1990) pp. 24-35.
- [79] S. Kramer, "Linus IV – An Experiment in Computer Security," *Proceedings of the 1984 Symposium on Security and Privacy* (Apr. 1984) pp. 24-31.
- [80] S. Kurzban, "Viruses and Worms -- What Can You Do?," *SIGSAC Review* **7**(1) pp. 16-32.
- [81] N. Lai and T. Gray, "Strengthening Discretionary Access Controls to Inhibit Trojan Horses and Computer Viruses," *1988 Summer USENIX Conference Proceedings* (June 1988) pp. 275-286.
- [82] B. Lampson, "Hints for Computer System Design," *IEEE Software* **1**(1) (Jan. 1984) pp. 11-28.
- [83] R. Levin, *Computer Virus Handbook*, McGraw-Hill Book Co., New York City, NY (1990).
- [84] T. Levin, S. Padilla, and C. Irvine, "A Formal Model for UNIX Setuid," *Proceedings of the 1989 Symposium on Security and Privacy* (May 1989) pp. 73-83.
- [85] P. Lewis, "The Executive Computer: A Virus Carries Fatal Complications," *New York Times* (June 26, 1988) p. C-11.
- [86] J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part III – Algorithms, Modes, and Identifiers*, RFC-1115 (Aug. 1989).
- [87] S. Lipner, "Non-Discretionary Controls for Commercial Applications," *Proceedings of the 1982 Symposium on Security and Privacy* (Apr. 1982) pp. 2-10.
- [88] T. Lunt and R. Jagannathan, "A Prototype Real-Time Intrusion-Detection Expert System," *Proceedings of the 1988 Symposium on Security and Privacy* (Apr. 1988) pp. 59-66.
- [89] J. McAfee, "4096 and 1260 Viruses (PC)," *Virus-L Digest* **3**(27) (Jan. 31, 1990), submitted by A. Roberts.

- [90] J. McAfee and C. Haynes, *Computer Viruses, Worms, Data Diddlers, Killer Programs, and Other Threats to Your System*, St. Martin's Press, New York City, NY (1989).
- [91] J. Markoff, "'Virus' in Military Computers Disrupts Systems Nationwide," *New York Times* (Nov. 4, 1988) p. A-1.
- [92] J. Markoff, "Top-Secret, And Vulnerable," *New York Times* (Apr. 25, 1988) p. A-1.
- [93] J. Markoff, "Student Says Error in Experiment Jammed a Network of Computers," *New York Times* (Jan. 19, 1990) p. A-19.
- [94] V. McLellan, "Computer Systems Under Siege," *New York Times* (Jan. 31, 1989) p. C-3.
- [95] R. Merkle, "A Fast Software One Way Hash Function," *unpublished*.
- [96] G. Miller, S. Sutton, M. Matthews, J. Yip, and T. Thomas, "Integrity Mechanisms in a Secure UNIX: GOULD UTX/32S," *AIAA/ASIS/DODCI Second Aerospace Computer Security Conference: A Collection of Technical Papers* (Dec. 1986) pp. 19-26.
- [97] W. Murray, "The Application of Epidemiology to Computer Viruses," *Computers and Security* 7(1) (Feb. 1988) pp. 139-150.
- [98] P. Neumann and D. Parker, "A Summary of Computer Misuse Techniques," *Twelfth National Computer Security Conference Proceedings* (Oct. 1989) pp. 396-407.
- [99] A. Norman, *Computer Insecurity*, Chapman and Hall, New York City, NY (1983).
- [100] J. Page, "An Assured Pipeline Integrity Scheme for Virus Protection," *Twelfth National Computer Security Conference Proceedings* (Oct. 1989) pp. 369-377.
- [101] D. Parker, *Crime by Computer*, Charles Scribner's Sons, New York City, NY (1976).
- [102] J. Picciotto, "The Design of an Effective Auditing Subsystem," *Proceedings of the 1987 Symposium on Security and Privacy* (Apr. 1987) pp. 13-22.
- [103] G. Popek and B. Walker, *The LOCUS Distributed System Architecture*, The MIT Press, Cambridge, MA (1985).
- [104] M. Pozzo and T. Gray, "A Model for the Containment of Computer Viruses," *AIAA/ASIS/DODCI Second Aerospace Computer Security Conference* (Dec. 1986) pp. 11-18.
- [105] M. Pozzo and T. Gray, "An Approach to Containing Computer Viruses," *Computers and Security* 6(4) (Aug. 1987) pp. 321-331.
- [106] B. Randell, P. Lee, and P. Treleaven, "Reliability Issues in Computing System Design,"

- Computing Surveys* **10**(2) (June 1978) pp. 167-196.
- [107] D. Ritchie, "Joy of Reproduction," USENET newsgroup *net.lang.c* (Nov. 4, 1982).
- [108] R. Roberts, *Computer Viruses*, Compute! Books, Greensboro, NC (1988).
- [109] J. Rochlis and M. Eichen, "With Microscope and Tweezers: The Worm from MIT's Perspective," *CACM* **32**(6) (June 1989) pp. 689-698.
- [110] J. Saltzer and M. Schroeder, "The Protection of Information in Computer Systems," *Proceedings of the IEEE* **63**(9) (Sep. 1975) pp. 1278-1308.
- [111] P. Samuelson, "Can Hackers Be Sued for Damages Caused by Computer Viruses?," *CACM* **32**(6) (June 1989) pp. 666-669.
- [112] O. Saydjari, J. Beckman, and J. Leaman, "Locking Computers Securely," *Tenth National Computer Security Conference Proceedings* (Sep. 1987) pp. 129-141.
- [113] O. Saydjari, J. Beckman, and J. Leaman, "LOCK Trek: Navigating Uncharted Space," *Proceedings of the 1989 Symposium on Security and Privacy* (May 1989) pp. 167-175.
- [114] R. Schatz, "New 'Virus' Infects NASA Macintoshes," *Washington Post* (Apr. 18, 1988), Washington Business section, p. 25.
- [115] J. Schoch and J. Hupp, "The 'Worm' Programs – Early Experiences with a Distributed Computation," *CACM* **25**(3) (Mar. 1982) pp. 172-180.
- [116] P. Scott, "Re: Faking Internet Mail [Re: RISKS-8.27]," *Forum on the Risks to the Public in Computers and Related Systems* **8**(28) (Feb. 19, 1989).
- [117] D. Seeley, "Password Cracking: A Game of Wits," *CACM* **32**(6) (June 1989) pp. 700-703.
- [118] D. Seeley, "A Tour of the Worm," *Proceedings of USENIX Winter '89* (Jan. 1989) pp. 287-304.
- [119] P. Singer, "Trying to Put a Brake on Computer Theft," *New York Times* (Mar. 2, 1986) p. WC-17.
- [120] K. Smith, "Tales of the Damned," *UNIX Review* **6**(2) (Feb. 1988) pp. 45-50.
- [121] T. Smith, "User Definable Domains as a Mechanism for Implementing the Least Privilege Principle," *Ninth National Computer Security Conference Proceedings* (Sep. 1986) pp. 143-148.
- [122] E. Spafford, "Crisis and Aftermath," *CACM* **32**(6) (June 1989) pp. 678-687.

- [123] E. Spafford, "The Internet Worm Program: An Analysis," *ACM Computer Communications Review* **19**(1) (Jan. 1989).
- [124] E. Spafford, K. Heaphy, and D. Ferbrache, *Computer Viruses: Dealing with Electronic Vandalism and Programmed Threats*, ADAPSO, Arlington, VA (1989).
- [125] C. Stoll, "An Epidemiology of Viruses & Network Worms," *Twelfth National Computer Security Conference Proceedings* (Oct. 1989) pp. 369-377.
- [126] H. Teng, K. Chen, and S. Lu, "Adaptive Real-Time Anomaly Detection Using Inductively Generated Sequential Patterns," *Proceedings of the 1990 Symposium on Research in Security and Privacy* (May 1990) pp. 278-284.
- [127] K. Thompson, "Reflections on Trusting Trust," *Communications of the ACM* **27**(8) (Aug. 1984) pp. 761-763.
- [128] M. Todd, "Man Catches Computer Virus!," *Weekly World News* p. 29 (June 18, 1991).
- [129] United States Comptroller General, "Computer-Related Crimes in Federal Programs," Report FGMSD-76-27, United States Government Printing Office, Washington, D. C. (Apr. 27, 1976).
- [130] United States Congress Office of Technology Assessment, *Defending Secrets, Sharing Data: New Locks and Keys for Electronic Information*, Report OTA-CIT-310, United States Government Printing Office, Washington, D. C. (Oct. 1987).
- [131] Virgil, *The Aeneid*, Random House, New York City, NY (1983).
- [132] C. von Rospach, "How to Post a Fake," *Forum on the Risks to the Public in Computers and Related Systems* **4**(75) (Apr. 20, 1987).
- [133] V. Voydock and S. Kent, "Security Mechanisms in High-Level Network Protocols," *Computing Surveys* **15**(2) (June 1983) pp. 135-171.
- [134] S. Wiseman, "Preventing Viruses in Computer Systems," *Computers and Security* **8**(5) (Aug. 1989) pp. 427-432.
- [135] I. Witten, "Computer (in)security: Infiltrating Open Systems," *Abacus* **4**(4) (1987) pp. 7-25.
- [136] P. Wood and S. Kochan, *UNIX™ System Security*, Hayden Books, Indianapolis, IN (1985).
- [137] C. Young, "Taxonomy of Computer Virus Defense Mechanisms," *Tenth National Computer Security Conference Proceedings* (Sep. 1987) pp. 220-225.

## **Sidebar 1 – The First Trojan Horse**

*There are many contradictory versions of this story; it appears only briefly in The Odyssey ([68], Book VIII), but later writers elaborated it considerably. Aeneas, a Trojan survivor of the sacking of the city, told the following version to Queen Dido of Carthage during his wanderings that ended with the founding of Rome ([131], Book II).*

After many years of besieging Troy and failing to take the city, the Greeks, on the advice of Athene, their patron goddess, built a large wooden horse in which many Greek soldiers hid. The horse was inscribed with a prayer to Athene to grant the Greeks safe passage home, and then the Greek army left.

The next morning, the Trojans discovered the siege had been lifted and went to examine the wooden horse. One of the elders, Thymoetes, noticed the inscription, and urged the horse be brought into the city and placed in Athene's temple. Others counseled that the horse must be destroyed; Laocoon, a priest of Apollo, threw a spear against the horse's belly as he cried that he did not trust Greeks bearing gifts.

Meanwhile, shepherds allied with the Trojans brought over a Greek soldier named Sinon. Sinon explained that the Greeks had desecrated Apollo's shrine and killed a virgin attendant in a raid, so to appease Apollo they had to sacrifice one of their men. Sinon was chosen. He promptly fled and was abandoned when the Greeks left for home. As for the horse, Sinon claimed that one night Odysseus and Diomedes desecrated Athene's shrine, turning their protecting goddess against them. Calchas, the Greeks' priest, advised that the horse must be built to appease the goddess before they could leave; and the horse was made so big to keep the Trojans from moving it into their city, for if they did their triumph over the Greeks would be assured.

At that moment, two sea serpents slithered out of the waters and crushed Laocoon and his sons to death. Believing this to be retribution for his profaning an offering to Athene, the Trojans immediately breached the walls of the city and pulled the horse inside.

That night, as the Trojans celebrated, they did not notice Sinon slip out to the horse and open a trap door through which the Greek soldiers emerged, nor did they see the Greeks opening the gates to the city. The Greek forces had by this time returned, and they sacked the city. Aeneas and his companions alone escaped.

## Sidebar 2 – Anatomy of a Virus

This pseudocode fragment shows how a very simple computer virus works:

```
beginvirus:
  if spread-condition then begin
    for some set of target files do begin
      if target is not infected then begin
        determine where to place virus instructions
        copy instructions from beginvirus to endvirus
        into target
        alter target to execute added instructions
      end;
    end;
  end;
  perform some action
  goto beginning of infected program
endvirus:
```

First, the virus determines if it is to spread; if so, it locates a set of target files it is to infect, and copies itself into a convenient location within the target file. It then alters portions of the target to ensure the inserted code will be executed at some time. For example, the virus may append itself just beyond the end of the instruction space and then adjust the entry points used by the loader so that the added instructions will execute when the target program is next run. This is the *infection phase*. It then performs some other action (the *execution phase*). Finally, it returns control to the program currently being run. Note that the execution phase can be null and the instructions still constitute a virus; but if the infection phase is missing, the instructions are *not* a virus.

The Lehigh virus [62] had as a *spread-condition* that “there is an uninfected boot file on the disk;” the *set of target files* was “the uninfected boot file,” and *perform some action* was to increment a counter and test to see if the counter had reached 4; if so, it would erase the disk.

### Sidebar 3 – A Starting Point for Suggested Guidelines for UNIX-based Systems

This list of suggestions, intended as a starting point for a basic, “vanilla” UNIX-based computer system, may help prevent the introduction of malicious logic, like computer viruses, into the computer system, and also lessen the chances of accidentally invoking programs with that type of logic. Attackers can render these methods ineffective because the weaknesses they seek to patch are fundamental to the design and use of the computer system, and anything effective would require changing the system more than is practical. Still, following these suggestions may help.

More details on UNIX security in general may be found in [33], [50], [53], and [136].

1. Set the environment variables (such as **PATH**) to access trusted programs before accessing untrusted programs of the same name.

The UNIX shell checks the value of the variable **PATH** for a list of directories to check for programs. The system administrator had put the current working directory before the system directories in the example in §6.1., Hence the user’s directory listing program, not the system one, was executed.

2. Do not execute a program obtained from an untrusted source without checking the source code thoroughly.

This rule presumes that the underlying computing base (compiler, loader, operating system, etc.) are all uncorrupted; if this assumption is false, malicious logic may be inserted during compilation, linking, or execution. An obvious corollary is to test all such software in an environment with very limited privileges before installing it, and *never* to test the program where it can access critical or irreplaceable files, or as a highly-privileged user.

3. Design and implement some auditing scheme to ensure that files’ access control permissions match the settings specified in an access control plan.

This requires first, that some security policy designating who has access to what files and how be created; and second, that some enforcement mechanism be implemented. Note the  *caveat*: if the audit log created by that mechanism, or the mechanism itself, can be tampered with, the introduction of malicious logic into the system can be done undetectably. However, depending on the security mechanisms implementing the auditing and the access to the log, this may require some sophistication. (Or, it may not.)

4. Check the integrity of system files to ensure they have not changed unexpectedly.

This is really a corollary to the previous rule. Note that the checksums computed at installation must be protected, since an attacker could change a file, then compute its new checksum and replace the stored checksum with it. Again, this requires that the underlying system be trusted to provide such protection to the checksum program, the stored checksums, and the audit program comparing the two.

5. Backups should be made regularly and kept as long as reasonable.

Typically, sites make both daily and weekly incremental backups (which save all files that have changed since the last incremental backup of the same period); then once a month they simply make a copy of all file systems. Enough of each kind is saved to be able to restore the

system to its current state. Notice that if restoring to eliminate a malicious program, the restored version of the program should also be *thoroughly* checked.

6. Discuss with your systems staff and users the reasons for, and effects of, any actions taken for security reasons.

The system staff should cultivate good relations with the users and vendors, should be certain to explain the reasons for all security policies, and should assist users whenever possible in providing a pleasant and secure working environment, acting as an intermediary between them and the vendors if need be. Users and staff should know what constitutes a breach of security, and there should be a well-designed set of procedures for handling breaches. Thinking through the best procedures for a particular installation carefully, putting them into place tactfully, and explaining them fully, will do far more to prevent security problems than any quick action.

7. All installations should keep the original distribution of the computer system in a safe place, and make and protect backups as well.

If malicious programs are determined to be rampant on the system, the administrators should reload the original compilation and installation software from the distribution medium and recompile and regenerate all system files after checking all sources thoroughly. This assumes that the (distributed) compilation and installation software is not infected and the program loading that software does not infect it. As always, the elements of trust are present here.

8. When reading backups, mount the backup medium in such a way that it cannot be changed or erased.

The reason is explained in the text. Note this means preventing modification access *by the hardware*, for example by removing the write ring from a tape. If the prevention mechanism is done in software, it can be infected and/or disabled by a malicious program. Here, the element of trust is in the hardware mechanism working correctly.

9. Access privileged accounts only when necessary, and then for as brief a time as possible.

Should someone using a privileged account accidentally execute a program containing a computer virus, the virus will spread throughout the system rapidly. This is less likely to happen if those accounts are used only when necessary; even so, a window of vulnerability still exists. Computers designed with security in mind typically limit the power of privileged accounts, in some cases very drastically.

10. Write as few privileged programs as possible.

The more programs that can cross protection domain boundaries while executing, the more potential targets for the addition of malicious logic exist. This suggestion essentially recommends minimizing the number of programs that can be modified to provide an attacker with entry to the privileged state.

11. Do not use a smart terminal to access a privileged account.

12. If a smart terminal must be used to access a privileged account, never allow an inter-terminal communications program to write to the terminal, never read electronic mail from that terminal, and do not look at files the contents of which are unknown or suspect.

Note that the second version is much weaker, because a malicious program could tamper with an executable program and cause it to display the control sequences to produce the req-

uisite commands from the terminal. The privileged user executing such a command springs the trap. Any file the malicious program could write to can be similarly booby-trapped.

13. Prevent users from accessing devices and memory directly.

If memory and devices are objects addressable by the user, the access control plan described earlier should include these objects and prevent direct access to them. Specifically, the device and memory files on UNIX systems should *never* have any *world* permissions set; this gives users direct access to memory and to the raw device, and allows them to bypass the UNIX access control mechanisms.

#### Sidebar 4 – Forums that Discuss Viruses

The VIRUS-L mailing list, moderated by Kenneth R. van Wyk, is a forum for discussing all aspects of computer viruses, especially existing computer viruses and countermeasures as well as theory. To subscribe, send an electronic mail message containing only the line

SUB VIRUS-L *your name*

to *LISTSERV@LEHIIBM1.BITNET*. Back issues of the digest are available by anonymous ftp from *IBM1.CC.LEHIGH.EDU* or *cert.sei.cmu.edu*; users not on the internet may send to the above address an electronic mail message containing only the line

GET VIRUS-L LOG $yy$  $mm$  $x$

where  $yy$  is the last two digits of the year,  $mm$  the number of the month, and  $x$  a letter indicating the number of the week in the month. For example, LOG8901B refers to the digests issued in the second week of January, 1989.

The mailing list VALERT-L is used *only* to announce viruses; any discussion is relegated to VIRUS-L. To subscribe, send an electronic mail message containing only the line

SUB VALERT-L *your name*

to the above address. Messages sent to VALERT-L appear in the next VIRUS-L digest as well.

Peter Neumann of SRI International moderates the Forum on Risks to the Public in Computers and Related Systems, or RISKS, list. This mailing list focuses on the risks involved in computer technology, and has discussed implications of viruses, although with a thrust different than the VIRUS-L mailing list. To subscribe, if on the Internet, send an electronic mail message to *RISKS-request@CSL.SRI.COM*; if on BITNET, send an electronic mail message containing only the line

SUBSCRIBE MD4H *your name*

to *LISTSERV@CMUCCVMA.BITNET*, or

SUBSCRIBE RISKS *your name*

to *LISTSERV@UGA.BITNET*, *LISTSERV@UBVM.BITNET*, or *LISTSERV@FINHUTC.BITNET*. Back issues of the digest are available by anonymous ftp from *crvax.sri.com* in the directory "RISKS:" and are named RISKS- $v.nn$  where  $v$  is the volume and  $nn$  the number within the volume.