

# **Secure Programming for Linux and Unix HOWTO**

**David A. Wheeler**

# **Secure Programming for Linux and Unix HOWTO**

by David A. Wheeler

v2.60 Edition

Published v2.60, 10 November 2000

Copyright © 1999, 2000 by David A. Wheeler

This paper provides a set of design and implementation guidelines for writing secure programs for Linux and Unix systems. Such programs include application programs used as viewers of remote data, web applications (including CGI scripts), network servers, and setuid/setgid programs. Specific guidelines for C, C++, Java, Perl, Python, TCL, and Ada95 are included.

This document is Copyright (C) 1999-2000 David A. Wheeler. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License (GFDL), Version 1.1 or any later version published by the Free Software Foundation; with the invariant sections being "About the Author", with no Front-Cover Texts, and no Back-Cover texts. A copy of the license is included in the section entitled "GNU Free Documentation License". This document is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

# Table of Contents

<b>1. Introduction</b> .....	<b>8</b>
<b>2. Background</b> .....	<b>11</b>
2.1. History of Unix, Linux, and Open Source / Free Software .....	11
2.1.1. Unix.....	11
2.1.2. Free Software Foundation .....	12
2.1.3. Linux .....	12
2.1.4. Open Source / Free Software .....	13
2.1.5. Comparing Linux and Unix .....	14
2.2. Security Principles .....	14
2.3. Types of Secure Programs .....	17
2.4. Paranoia is a Virtue .....	19
2.5. Why Did I Write This Document?.....	19
2.6. Sources of Design and Implementation Guidelines.....	20
2.7. Other Sources of Security Information .....	23
2.8. Document Conventions .....	23
<b>3. Summary of Linux and Unix Security Features</b> .....	<b>25</b>
3.1. Processes .....	26
3.1.1. Process Attributes .....	27
3.1.2. POSIX Capabilities .....	28
3.1.3. Process Creation and Manipulation .....	29
3.2. Files.....	30
3.2.1. Filesystem Object Attributes.....	30
3.2.2. Creation Time Initial Values .....	33
3.2.3. Changing Access Control Attributes .....	34
3.2.4. Using Access Control Attributes .....	34
3.2.5. Filesystem Hierarchy .....	35
3.3. System V IPC.....	35
3.4. Sockets and Network Connections .....	36
3.5. Signals.....	38
3.6. Quotas and Limits .....	39
3.7. Dynamically Linked Libraries .....	39

3.8. Audit .....	41
3.9. PAM .....	41
<b>4. Validate All Input.....</b>	<b>43</b>
4.1. Command line .....	45
4.2. Environment Variables .....	45
4.2.1. Some Environment Variables are Dangerous.....	46
4.2.2. Environment Variable Storage Format is Dangerous.....	46
4.2.3. The Solution - Extract and Erase .....	47
4.3. File Descriptors .....	49
4.4. File Contents .....	49
4.5. Web-Based Applications (Especially CGI Scripts) .....	50
4.6. Other Inputs .....	50
4.7. Human Language (Locale) Selection.....	51
4.7.1. How Locales are Selected .....	51
4.7.2. Locale Support Mechanisms.....	52
4.7.3. Legal Values.....	53
4.7.4. Bottom Line .....	54
4.8. Character Encoding.....	54
4.9. Limit Valid Input Time and Load Level .....	58
<b>5. Avoid Buffer Overflow .....</b>	<b>59</b>
5.1. Dangers in C/C++ .....	59
5.2. Library Solutions in C/C++ .....	60
5.2.1. Standard C Library Solution .....	61
5.2.2. Static and Dynamically Allocated Buffers.....	62
5.2.3. strncpy and strlcat .....	64
5.2.4. libmib .....	65
5.2.5. Libsafe.....	65
5.2.6. Other Libraries .....	67
5.3. Compilation Solutions in C/C++ .....	67
5.4. Other Languages .....	68
<b>6. Structure Program Internals and Approach .....</b>	<b>70</b>
6.1. Secure the Interface.....	70
6.2. Minimize Privileges .....	70

6.2.1. Minimize the Privileges Granted .....	71
6.2.2. Minimize the Time the Privilege Can Be Used .....	73
6.2.3. Minimize the Time the Privilege is Active .....	74
6.2.4. Minimize the Modules Granted the Privilege .....	74
6.2.5. Consider Using FSUID To Limit Privileges .....	75
6.2.6. Consider Using Chroot to Minimize Available Files .....	75
6.2.7. Consider Minimizing the Accessible Data .....	77
6.3. Avoid Creating Setuid/Setgid Scripts .....	78
6.4. Configure Safely and Use Safe Defaults .....	78
6.5. Fail Safe .....	79
6.6. Avoid Race Conditions .....	80
6.6.1. Sequencing (Non-Atomic) Problems .....	81
6.6.2. Locking .....	83
6.6.2.1. Using Files as Locks .....	83
6.6.2.2. Other Approaches to Locking .....	85
6.7. Trust Only Trustworthy Channels .....	86
6.8. Use Internal Consistency-Checking Code .....	88
6.9. Self-limit Resources .....	88
<b>7. Carefully Call Out to Other Resources .....</b>	<b>89</b>
7.1. Limit Call-outs to Valid Values .....	89
7.2. Call Only Interfaces Intended for Programmers .....	92
7.3. Check All System Call Returns .....	92
7.4. Avoid Using vfork(2) .....	93
7.5. Counter Web Bugs When Retrieving Embedded Content .....	94
<b>8. Send Information Back Judiciously .....</b>	<b>96</b>
8.1. Minimize Feedback .....	96
8.2. Handle Full/Unresponsive Output .....	96
8.3. Control Data Formatting .....	97
<b>9. Language-Specific Issues .....</b>	<b>100</b>
9.1. C/C++ .....	100
9.2. Perl .....	102
9.3. Python .....	103
9.4. Shell Scripting Languages (sh and csh Derivatives) .....	104

9.5. Ada.....	105
9.6. Java.....	105
9.7. TCL.....	110
<b>10. Special Topics .....</b>	<b>112</b>
10.1. Passwords.....	112
10.2. Random Numbers .....	113
10.3. Specially Protect Secrets (Passwords and Keys) in User Memory .....	113
10.4. Cryptographic Algorithms and Protocols .....	114
10.5. PAM .....	116
10.6. Tools.....	117
10.7. Miscellaneous .....	118
<b>11. Conclusion .....</b>	<b>121</b>
<b>12. Bibliography .....</b>	<b>122</b>
<b>A. History .....</b>	<b>130</b>
<b>B. Acknowledgements.....</b>	<b>131</b>
<b>C. About the Documentation License.....</b>	<b>133</b>
<b>D. GNU Free Documentation License.....</b>	<b>135</b>
<b>E. Endorsements .....</b>	<b>145</b>
<b>F. About the Author .....</b>	<b>146</b>

# List of Tables

4-1. Illegal UTF-8 initial sequences.....	56
---	----

# Chapter 1. Introduction

*A wise man attacks the city of the  
mighty and pulls down the  
stronghold in which they trust.*

*Proverbs 21:22 (NIV)*

This paper describes a set of design and implementation guidelines for writing secure programs on Linux and Unix systems. For purposes of this paper, a “secure program” is a program that sits on a security boundary, taking input from a source that does not have the same access rights as the program. Such programs include application programs used as viewers of remote data, web applications (including CGI scripts), network servers, and setuid/setgid programs. This paper does not address modifying the operating system kernel itself, although many of the principles discussed here do apply. These guidelines were developed as a survey of “lessons learned” from various sources on how to create such programs (along with additional observations by the author), reorganized into a set of larger principles. This paper includes specific guidance for a number of languages, including C, C++, Java, Perl, Python, TCL, and Ada95.

This paper does not cover assurance measures, software engineering processes, and quality assurance approaches, which are important but widely discussed elsewhere. Such measures include testing, peer review, configuration management, and formal methods. Documents specifically identifying sets of development assurance measures for security issues include the Common Criteria [CC 1999] and the System Security Engineering Capability Maturity Model [SSE-CMM 1999]. More general sets of software engineering methods or processes are defined in documents such as the Software Engineering Institute’s Capability Maturity Model for Software (SE-CMM), ISO 9000 (along with ISO 9001 and ISO 9001-3), and ISO 12207.

This paper does not discuss how to configure a system (or network) to be secure in a given environment. This is clearly necessary for secure use of a given program, but a great many other documents discuss secure configurations. An excellent general book on configuring Unix-like systems to be secure is Garfinkel [1996]. Other books for securing Unix-like systems include Anonymous [1998]. You can also find information



on configuring Unix-like systems at web sites such as <http://www.unixtools.com/security.html>. Information on configuring a Linux system to be secure is available in a wide variety of documents including Fenzi [1999], Seifried [1999], Wreski [1998], and Anonymous [1999]. For Linux systems (and eventually other Unix-like systems), you may want to examine the Bastille Hardening System, which attempts to “harden” or “tighten” the Linux operating system. You can learn more about Bastille at <http://www.bastille-linux.org>; it is available for free under the General Public License (GPL).

This paper assumes that the reader understands computer security issues in general, the general security model of Unix-like systems, and the C programming language. This paper does include some information about the Linux and Unix programming model for security.

This paper covers all Unix-like systems, including Linux and the various strains of Unix, and it particularly stresses Linux and provides details about Linux specifically. There are several reasons for this, but a simple reason is popularity. According to a 1999 survey by IDC, significantly more servers (counting both Internet and intranet servers) were installed in 1999 with Linux than with all Unix operating system types combined (25% for Linux versus 15% for all Unix system types combined; note that Windows NT came in with 38% compared to the 40% of all Unix-like servers) [Shankland 2000]. A survey by Zoebelin in April 1999 found that, of the total number of servers deployed on the Internet in 1999 (running at least ftp, news, or http (WWW)), the majority were running Linux (28.5%), with others trailing (24.4% for all Windows 95/98/NT combined, 17.7% for Solaris or SunOS, 15% for the BSD family, and 5.3% for IRIX). Advocates will notice that the majority of servers on the Internet (around 66%) were running Unix-like systems, while only around 24% ran a Microsoft Windows variant. Finally, the original version of this document only discussed Linux, so although its scope has expanded, the Linux information is still noticeably dominant. If you know relevant information not already included here, please let me know.

You can find the master copy of this document at <http://www.dwheeler.com/secure-programs>. This document is also part of the Linux Documentation Project (LDP) at <http://www.linuxdoc.org>. It's also mirrored in several other places. Please note that these mirrors, including the LDP copy and/or the copy in your distribution, may be older than the master copy. I'd like to hear comments on this

document, but please do not send comments until you've checked to make sure that your comment is valid for the latest version.

This document is (C) 1999-2000 David A. Wheeler and is covered by the GNU Free Documentation License (GFDL); see Appendix C and Appendix D for more information.

Chapter 2 discusses the background of Unix, Linux, and security. Chapter 3 describes the general Unix and Linux security model, giving an overview of the security attributes and operations of processes, filesystem objects, and so on. This is followed by the meat of this paper, a set of design and implementation guidelines for developing applications on Linux and Unix systems. The paper ends with conclusions in Chapter 11, followed by a lengthy bibliography and appendices.

The design and implementation guidelines are divided into categories which I believe emphasize the programmer's viewpoint. Programs accept inputs, process data, call out to other resources, and produce output; notionally all security guidelines fit into one of these categories. I've divided processing data into further categories: avoiding buffer overflows (which in some cases can also be considered an input issue), structuring program internals and approach, language-specific information, and special topics. The actual chapter layout was reordered slightly to be easier to follow. Thus, the document chapters on guidelines discuss validating all input (Chapter 4), avoiding buffer overflows (Chapter 5), structuring program internals and approach (Chapter 6), carefully calling out to other resources (Chapter 7), judiciously sending information back (Chapter 8), language-specific information (Chapter 9), and finally information on special topics such as how to acquire random numbers (Chapter 10).

# Chapter 2. Background

*I issued an order and a search was made, and it was found that this city has a long history of revolt against kings and has been a place of rebellion and sedition.*

*Ezra 4:19 (NIV)*

## 2.1. History of Unix, Linux, and Open Source / Free Software

### 2.1.1. Unix

In 1969-1970, Kenneth Thompson, Dennis Ritchie, and others at AT&T Bell Labs began developing a small operating system on a little-used PDP-7. The operating system was soon christened Unix, a pun on an earlier operating system project called MULTICS. In 1972-1973 the system was rewritten in the programming language C, an unusual step that was visionary: due to this decision, Unix was the first widely-used operating system that could switch from and outlive its original hardware. Other innovations were added to Unix as well, in part due to synergies between Bell Labs and the academic community. In 1979, the “seventh edition” (V7) version of Unix was released, the grandfather of all extant Unix systems.

After this point, the history of Unix becomes somewhat convoluted. The academic community, led by Berkeley, developed a variant called the Berkeley Software Distribution (BSD), while AT&T continued developing Unix under the names “System III” and later “System V”. In the late 1980’s through early 1990’s the “wars” between these two major strains raged. After many years each variant adopted many of the key features of the other. Commercially, System V won the “standards wars” (getting most

of its interfaces into the formal standards), and most hardware vendors switched to AT&T's System V. However, System V ended up incorporating many BSD innovations, so the resulting system was more a merger of the two branches. The BSD branch did not die, but instead became widely used for research, for PC hardware, and for single-purpose servers (e.g., many web sites use a BSD derivative).

The result was many different versions of Unix, all based on the original seventh edition. Most versions of Unix were proprietary and maintained by their respective hardware vendor, for example, Sun Solaris is a variant of System V. Three versions of the BSD branch of Unix ended up as open source: FreeBSD (concentrating on ease-of-installation for PC-type hardware), NetBSD (concentrating on many different CPU architectures), and a variant of NetBSD, OpenBSD (concentrating on security). More general information can be found at

<http://www.datametrics.com/tech/unix/uxhistry/brf-hist.htm>. Much more information about the BSD history can be found in [McKusick 1999] and <ftp://ftp.freebsd.org/pub/FreeBSD/FreeBSD-current/src/share/misc/bsd-family-tree>.

Those interested in reading an advocacy piece that presents arguments for using Unix-like systems should see <http://www.unix-vs-nt.org>.

## 2.1.2. Free Software Foundation

In 1984 Richard Stallman's Free Software Foundation (FSF) began the GNU project, a project to create a free version of the Unix operating system. By free, Stallman meant software that could be freely used, read, modified, and redistributed. The FSF successfully built a vast number of useful components, including a C compiler (gcc), an impressive text editor (emacs), and a host of fundamental tools. However, in the 1990's the FSF was having trouble developing the operating system kernel [FSF 1998]; without a kernel the rest of their software would not work.

## 2.1.3. Linux

In 1991 Linus Torvalds began developing an operating system kernel, which he named

“Linux” [Torvalds 1999]. This kernel could be combined with the FSF material and other components (in particular some of the BSD components and MIT’s X-windows software) to produce a freely-modifiable and very useful operating system. This paper will term the kernel itself the “Linux kernel” and an entire combination as “Linux”. Note that many use the term “GNU/Linux” instead for this combination.

In the Linux community, different organizations have combined the available components differently. Each combination is called a “distribution”, and the organizations that develop distributions are called “distributors”. Common distributions include Red Hat, Mandrake, SuSE, Caldera, Corel, and Debian. There are differences between the various distributions, but all distributions are based on the same foundation: the Linux kernel and the GNU glibc libraries. Since both are covered by “copyleft” style licenses, changes to these foundations generally must be made available to all, a unifying force between the Linux distributions at their foundation that does not exist between the BSD and AT&T-derived Unix systems. This paper is not specific to any Linux distribution; when it discusses Linux it presumes Linux kernel version 2.2 or greater and the C library glibc 2.1 or greater, valid assumptions for essentially all current major Linux distributions.

#### **2.1.4. Open Source / Free Software**

Increased interest in software that is freely shared has made it increasingly necessary to define and explain it. A widely used term is “open source software”, which is further defined in [OSI 1999]. Eric Raymond [1997, 1998] wrote several seminal articles examining its various development processes. Another widely-used term is “free software”, where the “free” is short for “freedom”: the usual explanation is “free speech, not free beer.” Neither phrase is perfect. The term “free software” is often confused with programs whose executables are given away at no charge, but whose source code cannot be viewed, modified, or redistributed. Conversely, the term “open source” is sometime (ab)used to mean software whose source code is visible, but for which there are limitations on use, modification, or redistribution. This paper uses the term “open source” for its usual meaning, that is, software which has its source code freely available for use, viewing, modification, and redistribution; a more detailed

definition is contained in the Open Source Definition (<http://www.opensource.org/osd.html>). In some cases, a difference in motive is suggested; those preferring the term “free software” wish to strongly emphasize the need for freedom, while those using the term may have other motives (e.g., higher reliability) or simply wish to appear less strident. For information on this definition of free software, and the motivations behind it, can be found at <http://www.fsf.org>.

Those interested in reading advocacy pieces for open source software and free software should see <http://www.opensource.org> and <http://www.fsf.org>. There are other papers which examine such software, for example, Miller [1995] found that the open source software were noticeably more reliable than proprietary software (using their measurement technique, which measured resistance to crashing due to random input).

## 2.1.5. Comparing Linux and Unix

This paper uses the term “Unix-like” to describe systems intentionally like Unix. In particular, the term “Unix-like” includes all major Unix variants and Linux distributions.

Linux is not derived from Unix source code, but its interfaces are intentionally like Unix. Therefore, Unix lessons learned generally apply to both, including information on security. Most of the information in this paper applies to any Unix-like system. Linux-specific information has been intentionally added to enable those using Linux to take advantage of Linux’s capabilities.

Unix-like systems share a number of security mechanisms, though there are subtle differences and not all systems have all mechanisms available. All include user and group ids (uids and gids) for each process and a filesystem with read, write, and execute permissions (for user, group, and other). See Thompson [1974] and Bach [1986] for general information on Unix systems, including their basic security mechanisms. Chapter 3 summarizes key security features of Unix and Linux.

## 2.2. Security Principles

There are many general security principles which you should be familiar with; consult a general text on computer security such as [Pfleeger 1997]. Often computer security goals are described in terms of three overall goals:

- *Confidentiality* (also known as secrecy), meaning that the computing system's assets are accessible only by authorized parties.
- *Integrity*, meaning that the assets can only be modified by authorized parties in authorized ways.
- *Availability*, meaning that the assets are accessible to the authorized parties in a timely manner (as determined by the systems requirements). The failure to meet this goal is called a denial of service.

Some people define additional security goals, while others lump those additional goals as special cases of these three goals. For example, some separately identify non-repudiation as a goal; this is the ability to “prove” that a sender sent or receiver received a message, even if the sender or receiver wishes to deny it later. Privacy is sometimes addressed separately from confidentiality; some define this as protecting the confidentiality of a *user* (e.g., their identity) instead of the data. Most goals require identification and authentication, which is sometimes listed as a separate goal. Often auditing (also called accountability) is identified as a desirable security goal. Sometimes “access control” and “authenticity” are listed separately as well. In any case, it is important to identify your program's overall security goals, no matter how you group those goals together, so that you'll know when you've met them.

Sometimes these goals are a response to a known set of threats, and sometimes some of these goals are required by law. For example, for U.S. banks and other financial institutions, there's a new privacy law called the “Gramm-Leach-Bliley” (GLB) Act. This law mandates disclosure of personal information shared and means of securing that data, requires disclosure of personal information that will be shared with third parties, and directs institutions to give customers a chance to opt out of data sharing. [Jones 2000]

Saltzer [1974] and later Saltzer and Schroeder [1975] list the following principles of the design of secure protection systems, which are still valid:

- *Least privilege.* Each user and program should operate using the fewest privileges possible. This principle limits the damage from an accident, error, or attack. It also reduces the number of potential interactions among privileged programs, so unintentional, unwanted, or improper uses of privilege are less likely to occur. This idea can be extended to the internals of a program: only the smallest portion of the program which needs those privileges should have them.
- *Economy of mechanism.* The protection system's design should be simple and small as possible. In their words, "techniques such as line-by-line inspection of software and physical examination of hardware that implements protection mechanisms are necessary. For such techniques to be successful, a small and simple design is essential."
- *Open design.* The protection mechanism must not depend on attacker ignorance. Instead, the mechanism should be public, depending on the secrecy of relatively few (and easily changeable) items like passwords or private keys. An open design makes extensive public scrutiny possible, and it also makes it possible for users to convince themselves that the system about to be used is adequate. Frankly, it isn't realistic to try to maintain secrecy for a system that is widely distributed; decompilers and subverted hardware can quickly expose any "secrets" in an implementation. Bruce Schneier argues that smart engineers should "demand open source code for anything related to security", as well as ensuring that it receives widespread review and that any identified problems are fixed [Schneier 1999].
- *Complete mediation.* Every access attempt must be checked; position the mechanism so it cannot be subverted. For example, in a client-server model, generally the server must do all access checking because users can build or modify their own clients.
- *Fail-safe defaults (e.g., permission-based approach).* The default should be denial of service, and the protection scheme should then identify conditions under which access is permitted.
- *Separation of privilege.* Ideally, access to objects should depend on more than one condition, so that defeating one protection system won't enable complete access.



- *Least common mechanism.* Minimize the amount and use of shared mechanisms (e.g. use of the /tmp or /var/tmp directories). Shared objects provide potentially dangerous channels for information flow and unintended interactions.
- *Psychological acceptability / Easy to use.* The human interface must be designed for ease of use so users will routinely and automatically use the protection mechanisms correctly. Mistakes will be reduced if the security mechanisms closely match the user's mental image of his or her protection goals.

## 2.3. Types of Secure Programs

Many different types of programs may need to be secure programs (as the term is defined in this paper). Some common types are:

- Application programs used as viewers of remote data. Programs used as viewers (such as word processors or file format viewers) are often asked to view data sent remotely by an untrusted user (this request may be automatically invoked by a web browser). Clearly, the untrusted user's input should not be allowed to cause the application to run arbitrary programs. It's usually unwise to support initialization macros (run when the data is displayed); if you must, then you must create a secure sandbox (a complex and error-prone task). Be careful of issues such as buffer overflow, discussed in Chapter 5, which might allow an untrusted user to force the viewer to run an arbitrary program.
- Application programs used by the administrator (root). Such programs shouldn't trust information that can be controlled by non-administrators.
- Local servers (also called daemons).
- Network-accessible servers (sometimes called network daemons).
- Web-based applications (including CGI scripts). These are a special case of network-accessible servers, but they're so common they deserve their own category.

Such programs are invoked indirectly via a web server, which filters out some attacks but nevertheless leaves many attacks that must be withstood.

- Applets (i.e., programs downloaded to the client for automatic execution). This is something Java is especially famous for, though other languages (such as Python) support mobile code as well. There are several security viewpoints here; the implementor of the applet infrastructure on the client side has to make sure that the only operations allowed are “safe” ones, and the writer of an applet has to deal with the problem of hostile hosts (in other words, you can’t normally trust the client). There is some research attempting to deal with running applets on hostile hosts, but frankly I’m sceptical of the value of these approaches and this subject is exotic enough that I don’t cover it further here.
- `setuid/setgid` programs. These programs are invoked by a local user and, when executed, are immediately granted the privileges of the program’s owner and/or owner’s group. In many ways these are the hardest programs to secure, because so many of their inputs are under the control of the untrusted user and some of those inputs are not obvious.

This paper merges the issues of these different types of program into a single set. The disadvantage of this approach is that some of the issues identified here don’t apply to all types of programs. In particular, `setuid/setgid` programs have many surprising inputs and several of the guidelines here only apply to them. However, things are not so clear-cut, because a particular program may cut across these boundaries (e.g., a CGI script may be `setuid` or `setgid`, or be configured in a way that has the same effect), and some programs are divided into several executables each of which can be considered a different “type” of program. The advantage of considering all of these program types together is that we can consider all issues without trying to apply an inappropriate category to a program. As will be seen, many of the principles apply to all programs that need to be secured.

There is a slight bias in this paper towards programs written in C, with some notes on other languages such as C++, Perl, Python, Ada95, and Java. This is because C is the most common language for implementing secure programs on Unix-like systems (other than CGI scripts, which tend to use Perl), and most other languages’ implementations call the C library. This is not to imply that C is somehow the “best” language for this

purpose, and most of the principles described here apply regardless of the programming language used.

## 2.4. Paranoia is a Virtue

The primary difficulty in writing secure programs is that writing them requires a different mindset, in short, a paranoid mindset. The reason is that the impact of errors (also called defects or bugs) can be profoundly different.

Normal non-secure programs have many errors. While these errors are undesirable, these errors usually involve rare or unlikely situations, and if a user should stumble upon one they will try to avoid using the tool that way in the future.

In secure programs, the situation is reversed. Certain users will intentionally search out and cause rare or unlikely situations, in the hope that such attacks will give them unwarranted privileges. As a result, when writing secure programs, paranoia is a virtue.

## 2.5. Why Did I Write This Document?

One question I've been asked is "why did you write this document"? Here's my answer: Over the last several years I've noticed that many developers for Linux and Unix seem to keep falling into the same security pitfalls, again and again. Auditors were slowly catching problems, but it would have been better if the problems weren't put into the code in the first place. I believe that part of the problem was that there wasn't a single, obvious place where developers could go and get information on how to avoid known pitfalls. The information was publicly available, but it was often hard to find, out-of-date, incomplete, or had other problems. Most such information didn't particularly discuss Linux at all, even though it was becoming widely used! That leads up to the answer: I developed this document in the hope that future software developers for Linux won't repeat past mistakes, resulting in an even more secure form of Linux. I added Unix, since it's often wise to make sure that programs can port between these systems. You can see a larger discussion of this at

[http://www.linuxsecurity.com/feature\\_stories/feature\\_story-6.html](http://www.linuxsecurity.com/feature_stories/feature_story-6.html).

A related question that could be asked is “why did you write your own document instead of just referring to other documents”? There are several answers:

- Much of this information was scattered about; placing the critical information in one organized document makes it easier to use.
- Some of this information is not written for the programmer, but is written for an administrator or user.
- Much of the available information emphasizes portable constructs (constructs that work on all Unix-like systems), and failed to discuss Linux at all. It’s often best to avoid Linux-unique abilities for portability’s sake, but sometimes the Linux-unique abilities can really aid security. Even if non-Linux portability is desired, you may want to support the Linux-unique abilities when running on Linux. And, by emphasizing Linux, I can include references to information that is helpful to someone targeting Linux that is not necessarily true for others.

## 2.6. Sources of Design and Implementation Guidelines

Several documents help describe how to write secure programs (or, alternatively, how to find security problems in existing programs), and were the basis for the guidelines highlighted in the rest of this paper.

For general-purpose servers and setuid/setgid programs, there are a number of valuable documents (though some are difficult to find without having a reference to them).

Matt Bishop [1996, 1997] has developed several extremely valuable papers and presentations on the topic, and in fact he has a web page dedicated to the topic at <http://olympus.cs.ucdavis.edu/~bishop/secprog.html>. AUSCERT has released a programming checklist [AUSCERT 1996] ([ftp://ftp.auscert.org.au/pub/auscert/papers/secure\\_programming\\_checklist](ftp://ftp.auscert.org.au/pub/auscert/papers/secure_programming_checklist)), based in

part on chapter 23 of Garfinkel and Spafford's book discussing how to write secure SUID and network programs [Garfinkel 1996] (<http://www.oreilly.com/catalog/puis>). Galvin [1998a] (<http://www.sunworld.com/swol-04-1998/swol-04-security.html>) described a simple process and checklist for developing secure programs; he later updated the checklist in Galvin [1998b] (<http://www.sunworld.com/sunworldonline/swol-08-1998/swol-08-security.html>). Sitaker [1999] (<http://www.pobox.com/~kragen/security-holes.html>) presents a list of issues for the "Linux security audit" team to search for. Shostack [1999] (<http://www.homeport.org/~adam/review.html>) defines another checklist for reviewing security-sensitive code. The NCSA [NCSA] (<http://www.ncsa.uiuc.edu/General/Grid/ACES/security/programming>) provides a set of terse but useful secure programming guidelines. Other useful information sources include the *Secure Unix Programming FAQ* [Al-Herbish 1999] (<http://www.whitefang.com/sup/>), the *Security-Audit's Frequently Asked Questions* [Graham 1999] (<http://lsap.org/faq.txt>), and Ranum [1998] (<http://www.clark.net/pub/mjr/pubs/pdf/>). Some recommendations must be taken with caution, for example, the BSD `setuid(7)` man page [Unknown] (<http://www.homeport.org/~adam/setuid.7.html>) recommends the use of `access(3)` without noting the dangerous race conditions that usually accompany it. Wood [1985] has some useful but dated advice in its "Security for Programmers" chapter. Bellovin [1994] (<http://www.research.att.com/~smb/talks>) includes useful guidelines and some specific examples, such as how to restructure an `ftpd` implementation to be simpler and more secure. FreeBSD [1999] (<http://www.freebsd.org/security/security.html>) [Quintero 1999] (<http://developer.gnome.org/doc/guides/programming-guidelines/book1.html>) is primarily concerned with GNOME programming guidelines, but it includes a section on security considerations. [Venema 1996] (<http://www.fish.com/security/murphy.html>) provides a detailed discussion (with examples) of some common errors when programming secure programs (widely-known or predictable passwords, burning yourself with malicious data, secrets in user-accessible data, and depending on other programs). [Sibert 1996] (<http://www.fish.com/security/maldata.html>) describes threats arising from malicious data.

There are many documents giving security guidelines for programs using the Common

Gateway Interface (CGI) to interface with the web. These include Van Biesbrouck [1996] (<http://www.csclub.uwaterloo.ca/u/mlvanbie/cgisec>), Gundavaram [unknown] (<http://language.perl.com/CPAN/doc/FAQs/cgi/perl-cgi-faq.html>), [Garfinkle 1997] (<http://webreview.com/wr/pub/97/08/08/bookshelf>) Kim [1996] (<http://www.eekim.com/pubs/cgibook>), Phillips [1995] (<http://www.go2net.com/people/paulp/cgi-security/safe-cgi.txt>), Stein [1999] (<http://www.w3.org/Security/Faq/www-security-faq.html>), [Peteanu 2000] (<http://members.home.net/razvan.peteanu>), and [Advosys 2000] (<http://advosys.ca/tips/web-security.html>).

There are many documents specific to a language, which are further discussed in the language-specific sections of this document. For example, the Perl distribution includes `perlsec(1)` (<http://www.perl.com/pub/doc/manual/html/pod/perlsec.html>), which describes how to use Perl more securely. The Secure Internet Programming site at <http://www.cs.princeton.edu/sip> is interested in computer security issues in general, but focuses on mobile code systems such as Java, ActiveX, and JavaScript; Ed Felten (one of its principles) co-wrote a book on securing Java ([McGraw 1999] ([www.securingjava.com](http://www.securingjava.com))) which is discussed in Section 9.6. Sun's security code guidelines provide some guidelines primarily for Java and C; it is available at <http://java.sun.com/security/seccodeguide.html>.

Yoder [1998] contains a collection of patterns to be used when dealing with application security. It's not really a specific set of guidelines, but a set of commonly-used patterns for programming that you may find useful. The Schmoo group maintains a web page linking to information on how to write secure code at <http://www.shmoo.com/securecode>

There are many documents describing the issue from the other direction (i.e., "how to crack a system"). One example is McClure [1999], and there's countless amounts of material from that vantage point on the Internet.

There's also a large body of information on vulnerabilities already identified in existing programs. This can be a useful set of examples of "what not to do," though it takes effort to extract more general guidelines from the large body of specific examples. There are mailing lists that discuss security issues; one of the most well-known is Bugtraq (<http://SecurityFocus.com/forums/bugtraq/faq.html>), which among other

things develops a list of vulnerabilities. The CERT Coordination Center (CERT/CC) is a major reporting center for Internet security problems which reports on vulnerabilities. The CERT/CC occasionally produces advisories that provide a description of a serious security problem and its impact, along with instructions on how to obtain a patch or details of a workaround; for more information see <http://www.cert.org>. Note that originally the CERT was a small computer emergency response team, but officially “CERT” doesn’t stand for anything now. The Department of Energy’s Computer Incident Advisory Capability (CIAC) (<http://ciac.llnl.gov/ciac>) also reports on vulnerabilities. These different groups may identify the same vulnerabilities but use different names. To resolve this problem, MITRE supports the Common Vulnerabilities and Exposures (CVE) list which creates a single unique identifier (“name”) for all publicly known vulnerabilities and security exposures identified by others; see <http://www.cve.mitre.org>. NIST’s ICAT is a searchable catalogue of computer vulnerabilities, taking the each CVE vulnerability and categorizing them so they can be searched and compared later; see <http://csrc.nist.gov/icat>.

This paper is a summary of what I believe are the most useful and important guidelines; my goal is a document that a good programmer can just read and then be fairly well prepared to implement a secure program. No single document can really meet this goal, but I believe the attempt is worthwhile. My goal is to strike a balance somewhere between a “complete list of all possible guidelines” (that would be unending and unreadable) and the various “short” lists available on-line that are nice and short but omit a large number of critical issues. When in doubt, I include the guidance; I believe in that case it’s better to make the information available to everyone in this “one stop shop” document. The organization presented here is my own (every list has its own, different structure), and some of the guidelines (especially the Linux-unique ones, such as those on capabilities and the fsuid value) are also my own. Reading all of the referenced documents listed above as well is highly recommended.

## 2.7. Other Sources of Security Information

Here are some other sources of security information: The Linux Security-Audit Project FAQ (<http://www.linuxhelp.org/lsap.shtml>) and Linux Kernel Auditing Project

(<http://www.lkap.org>) are dedicated to auditing Linux code for security issues.

## 2.8. Document Conventions

System manual pages are referenced in the format *name(number)*, where *number* is the section number of the manual. The pointer value that means “does not point anywhere” is called NULL; C compilers will convert the integer 0 to the value NULL in most circumstances where a pointer is needed, but note that nothing in the C standard requires that NULL actually be implemented by a series of all-zero bits. C and C++ treat the character ‘\0’ (ASCII 0) specially, and this value is referred to as NIL in this paper (this is usually called “NUL”, but “NUL” and “NULL” sound identical). Function and method names always use the correct case, even if that means that some sentences must begin with a lower case letter. I use the term “Unix-like” to mean Unix, Linux, or other systems whose underlying models are very similar to Unix; I can’t say POSIX, because there are systems such as Windows 2000 that implement portions of POSIX yet have vastly different security models.

An attacker is called an “attacker”, “cracker”, or “adversary”. Some journalists use the word “hacker” instead of “attacker”; this paper avoids this (mis)use, because many Linux and Unix developers refer to themselves as “hackers” in the traditional non-evil sense of the term. That is, to many Linux and Unix developers, the term “hacker” continues to mean simply an expert or enthusiast, particularly regarding computers.

This document uses the “new” or “logical” quoting system, instead of the traditional American quoting system: quoted information does not include any trailing punctuation if the punctuation is not part of the material being quoted. While this may cause a minor loss of typographical beauty, the traditional American system causes extraneous characters to be placed inside the quotes. These extraneous characters have no effect on prose but can be disastrous in code or computer commands. I use standard American (not British) spelling; I’ve yet to meet an English speaker on any continent who has trouble with this.



# Chapter 3. Summary of Linux and Unix Security Features

*Discretion will protect you, and  
understanding will guard you.*

*Proverbs 2:11 (NIV)*

Before discussing guidelines on how to use Linux or Unix security features, it's useful to know what those features are from a programmer's viewpoint. This section briefly describes those features that are widely available on nearly all Unix-like systems. However, note that there is considerable variation between different versions of Unix-like systems, and not all systems have the abilities described here. This chapter also notes some extensions or features specific to Linux; Linux distributions tend to be fairly similar to each other from the point-of-view of programming for security, because they all use essentially the same kernel and C library (and the GPL-based licenses encourage rapid dissemination of any innovations). This chapter doesn't discuss issues such as implementations of mandatory access control (MAC) which many Unix-like systems do not implement. If you already know what those features are, please feel free to skip this section.

Many programming guides skim briefly over the security-relevant portions of Linux or Unix and skip important information. In particular, they often discuss "how to use" something in general terms but gloss over the security attributes that affect their use. Conversely, there's a great deal of detailed information in the manual pages about individual functions, but the manual pages sometimes obscure key security issues with detailed discussions on how to use each individual function. This section tries to bridge that gap; it gives an overview of the security mechanisms in Linux that are likely to be used by a programmer, but concentrating specifically on the security ramifications. This section has more depth than the typical programming guides, focusing specifically on security-related matters, and points to references where you can get more details.

First, the basics. Linux and Unix are fundamentally divided into two parts: the kernel and "user space". Most programs execute in user space (on top of the kernel). Linux

supports the concept of “kernel modules”, which is simply the ability to dynamically load code into the kernel, but note that it still has this fundamental division. Some other systems (such as the HURD) are “microkernel” based systems; they have a small kernel with more limited functionality, and a set of “user” programs that implement the lower-level functions traditionally implemented by the kernel.

Some Unix-like systems have been extensively modified to support strong security, in particular to support U.S. Department of Defense requirements for Mandatory Access Control (level B1 or higher). This version of this paper doesn’t cover these systems or issues; I hope to expand to that in a future version.

When users log in, their usernames are mapped to integers marking their “UID” (for “user id”) and the “GID”s (for “group id”) that they are a member of. UID 0 is a special privileged user (role) traditionally called “root”; on most Unix-like systems (including Unix) root can overrule most security checks and is used to administrate the system. Processes are the only “subjects” in terms of security (that is, only processes are active objects). Processes can access various data objects, in particular filesystem objects (FSOs), System V Interprocess Communication (IPC) objects, and network ports. Processes can also set signals. Other security-relevant topics include quotas and limits, libraries, auditing, and PAM. The next few subsections detail this.

## 3.1. Processes

In Unix-like systems, user-level activities are implemented by running processes. Most Unix systems support a “thread” as a separate concept; threads share memory inside a process, and the system scheduler actually schedules threads. Linux does this differently (and in my opinion uses a better approach): there is no essential difference between a thread and a process. Instead, in Linux, when a process creates another process it can choose what resources are shared (e.g., memory can be shared). The Linux kernel then performs optimizations to get thread-level speeds; see `clone(2)` for more information. It’s worth noting that the Linux kernel developers tend to use the word “task”, not “thread” or “process”, but the external documentation tends to use the word process (so I’ll use that terminology here). When programming a multi-threaded application, it’s usually better to use one of the standard thread libraries that hide these

differences. Not only does this make threading more portable, but some libraries provide an additional level of indirection, by implementing more than one application-level thread as a single operating system thread; this can provide some improved performance on some systems for some applications.

### 3.1.1. Process Attributes

Here are typical attributes associated with each process in a Unix-like system:

- RUID, RGID - real UID and GID of the user on whose behalf the process is running
- EUID, EGID - effective UID and GID used for privilege checks (except for the filesystem)
- SUID, SGID - Saved UID and GID; used to support switching permissions “on and off” as discussed below. Not all Unix-like systems support this.
- supplemental groups - a list of groups (GIDs) in which this user has membership.
- umask - a set of bits determining the default access control settings when a new filesystem object is created; see `umask(2)`.
- scheduling parameters - each process has a scheduling policy, and those with the default policy `SCHED_OTHER` have the additional parameters `nice`, `priority`, and `counter`. See `sched_setscheduler(2)` for more information.
- limits - per-process resource limits (see below).
- filesystem root - the process’ idea of where the root filesystem begins; see `chroot(2)`.

Here are less-common attributes associated with processes:

- FSUID, FSGID - UID and GID used for filesystem access checks; this is usually equal to the EUID and EGID respectively. This is a Linux-unique attribute.
- capabilities - POSIX capability information; there are actually three sets of capabilities on a process: the effective, inheritable, and permitted capabilities. See

below for more information on POSIX capabilities. Linux kernel version 2.2 and greater support this; some other Unix-like systems do too, but it's not as widespread.

In Linux, if you really need to know exactly what attributes are associated with each process, the most definitive source is the Linux source code, in particular `/usr/include/linux/sched.h`'s definition of `task_struct`.

The portable way to create new processes is to use the `fork(2)` call. BSD introduced a variant called `vfork(2)` as an optimization technique. The bottom line with `vfork(2)` is simple: *don't* use it if you can avoid it. See Section 7.4 for more information.

Linux supports the Linux-unique `clone(2)` call. This call works like `fork(2)`, but allows specification of which resources should be shared (e.g., memory, file descriptors, etc.). Portable programs shouldn't use this call directly; as noted earlier, they should instead rely on threading libraries that use the call to implement threads.

This document is not a full tutorial on writing programs, so I will skip widely-available information handling processes. You can see the documentation for `wait(2)`, `exit(2)`, and so on for more information.

### 3.1.2. POSIX Capabilities

POSIX capabilities are sets of bits that permit splitting of the privileges typically held by root into a larger set of more specific privileges. POSIX capabilities are defined by a draft IEEE standard; they're not unique to Linux but they're not universally supported by other Unix-like systems either. Linux kernel 2.0 did not support POSIX capabilities, while version 2.2 added support for POSIX capabilities to processes. When Linux documentation (including this one) says "requires root privilege", in nearly all cases it really means "requires a capability" as documented in the capability documentation. If you need to know the specific capability required, look it up in the capability documentation.

In Linux, the eventual intent is to permit capabilities to be attached to files in the filesystem; as of this writing, however, this is not yet supported. There is support for transferring capabilities, but this is disabled by default. Linux version 2.2.11 added a

feature that makes capabilities more directly useful, called the “capability bounding set”. The capability bounding set is a list of capabilities that are allowed to be held by any process on the system (otherwise, only the special init process can hold it). If a capability does not appear in the bounding set, it may not be exercised by any process, no matter how privileged. This feature can be used to, for example, disable kernel module loading. A sample tool that takes advantage of this is LCAP at <http://pweb.netcom.com/~spoon/lcap/>.

More information about POSIX capabilities is available at <ftp://linux.kernel.org/pub/linux/libs/security/linux-privs>.

### 3.1.3. Process Creation and Manipulation

Processes may be created using `fork(2)`, the non-recommended `vfork(2)`, or the Linux-unique `clone(2)`; all of these system calls duplicate the existing process, creating two processes out of it. A process can execute a different program by calling `execve(2)`, or various front-ends to it (for example, see `exec(3)`, `system(3)`, and `popen(3)`).

When a program is executed, and its file has its `setuid` or `setgid` bit set, the process’ `EUID` or `EGID` (respectively) is usually set to the file’s value. This functionality was the source of an old Unix security weakness when used to support `setuid` or `setgid` scripts, due to a race condition. Between the time the kernel opens the file to see which interpreter to run, and when the (now-set-id) interpreter turns around and reopens the file to interpret it, an attacker might change the file (directly or via symbolic links).

Different Unix-like systems handle the security issue for `setuid` scripts in different ways. Some systems, such as Linux, completely ignore the `setuid` and `setgid` bits when executing scripts, which is clearly a safe approach. Most modern releases of SysVr4 and BSD 4.4 use a different approach to avoid the kernel race condition. On these systems, when the kernel passes the name of the set-id script to open to the interpreter, rather than using a pathname (which would permit the race condition) it instead passes the filename `/dev/fd/3`. This is a special file already opened on the script, so that there can be no race condition for attackers to exploit. Even on these systems I recommend against using the `setuid/setgid` shell scripts language for secure programs, as discussed below.

In some cases a process can affect the various UID and GID values; see `setuid(2)`, `seteuid(2)`, `setreuid(2)`, and the Linux-unique `setfsuid(2)`. In particular the saved user id (SUID) attribute is there to permit trusted programs to temporarily switch UIDs. Unix-like systems supporting the SUID use the following rules: If the RUID is changed, or the EUID is set to a value not equal to the RUID, the SUID is set to the new EUID. Unprivileged users can set their EUID from their SUID, the RUID to the EUID, and the EUID to the RUID.

The Linux-unique FSUID process attribute is intended to permit programs like the NFS server to limit themselves to only the filesystem rights of some given UID without giving that UID permission to send signals to the process. Whenever the EUID is changed, the FSUID is changed to the new EUID value; the FSUID value can be set separately using `setfsuid(2)`, a Linux-unique call. Note that non-root callers can only set FSUID to the current RUID, EUID, SEUID, or current FSUID values.

## 3.2. Files

On all Unix-like systems, the primary repository of information is the file tree, rooted at `/`. The file tree is a hierarchical set of directories, each of which may contain filesystem objects (FSOs).

In Linux, filesystem objects (FSOs) may be ordinary files, directories, symbolic links, named pipes (also called first-in first-outs or FIFOs), sockets (see below), character special (device) files, or block special (device) files (in Linux, this list is given in the `find(1)` command). Other Unix-like systems have an identical or similar list of FSO types.

Filesystem objects are collected on filesystems, which can be mounted and unmounted on directories in the file tree. A filesystem type (e.g., ext2 and FAT) is a specific set of conventions for arranging data on the disk to optimize speed, reliability, and so on; many people use the term “filesystem” as a synonym for the filesystem type.

### 3.2.1. Filesystem Object Attributes

Different Unix-like systems support different filesystem types. Filesystems may have slightly different sets of access control attributes and access controls can be affected by options selected at mount time. On Linux, the ext2 filesystem is currently the most popular filesystem, but Linux supports a vast number of filesystems. Most Unix-like systems tend to support multiple filesystems too.

Most filesystems on Unix-like systems store at least the following:

- owning UID and GID - identifies the “owner” of the filesystem object. Only the owner or root can change the access control attributes unless otherwise noted.
- permission bits - read, write, execute bits for each of user (owner), group, and other. For ordinary files, read, write, and execute have their typical meanings. In directories, the “read” permission is necessary to display a directory’s contents, while the “execute” permission is sometimes called “search” permission and is necessary to actually enter the directory to use its contents. In a directory “write” permission on a directory permits adding, removing, and renaming files in that directory; if you only want to permit adding, set the sticky bit noted below. Note that the permission values of symbolic links are never used; it’s only the values of their containing directories and the linked-to file that matter.
- “sticky” bit - when set on a directory, unlinks (removes) and renames of files in that directory are limited to the file owner, the directory owner, or root privileges. This is a very common Unix extension and is specified in the Open Group’s Single Unix Specification version 2. Old versions of Unix called this the “save program text” bit and used this to indicate executable files that should stay in memory. Systems that did this ensured that only root could set this bit (otherwise users could have crashed systems by forcing “everything” into memory). In Linux, this bit has no affect on ordinary files and ordinary users can modify this bit on the files they own: Linux’s virtual memory management makes this old use irrelevant.
- setuid, setgid - when set on an executable file, executing the file will set the process’ effective UID or effective GID to the value of the file’s owning UID or GID (respectively). All Unix-like systems support this. In Linux and System V systems,

when `setgid` is set on a file that does not have any execute privileges, this indicates a file that is subject to mandatory locking during access (if the filesystem is mounted to support mandatory locking); this overload of meaning surprises many and is not universal across Unix-like systems. In fact, the Open Group's Single Unix Specification version 2 for `chmod(3)` permits systems to ignore requests to turn on `setgid` for files that aren't executable if such a setting has no meaning. In Linux and Solaris, when `setgid` is set on a directory, files created in the directory will have their GID automatically reset to that of the directory's GID. The purpose of this approach is to support "project directories": users can save files into such specially-set directories and the group owner automatically changes. However, setting the `setgid` bit on directories is not specified by standards such as the Single Unix Specification [Open Group 1997].

- timestamps - access and modification times are stored for each filesystem object. However, the owner is allowed to set these values arbitrarily (see `touch(1)`), so be careful about trusting this information. All Unix-like systems support this.

The following are attributes are Linux-unique extensions on the `ext2` filesystem, though many other filesystems have similar functionality:

- immutable bit - no changes to the filesystem object are allowed; only root can set or clear this bit. This is only supported by `ext2` and is not portable across all Unix systems (or even all Linux filesystems).
- append-only bit - only appending to the filesystem object are allowed; only root can set or clear this bit. This is only supported by `ext2` and is not portable across all Unix systems (or even all Linux filesystems).

Other common extensions include some sort of bit indicating "cannot delete this file".

Many of these values can be influenced at mount time, so that, for example, certain bits can be treated as though they had a certain value (regardless of their values on the media). See `mount(1)` for more information about this. Some filesystems don't support some of these access control values; again, see `mount(1)` for how these filesystems are handled. In particular, many Unix-like systems support MS-DOS disks, which by default support very few of these attributes (and there's not standard way to define these



attributes). In that case, Unix-like systems emulate the standard attributes (possibly implementing them through special on-disk files), and these attributes are generally influenced by the `mount(1)` command.

It's important to note that, for adding and removing files, only the permission bits and owner of the file's *directory* really matter unless the Unix-like system supports more complex schemes (such as POSIX ACLs). Unless the system has other extensions, and stock Linux 2.2 doesn't, a file that has no permissions in its permission bits can still be removed if its containing directory permits it. Also, if an ancestor directory permits its children to be changed by some user or group, then any of that directory's descendents can be replaced by that user or group.

The draft IEEE POSIX standard on security defines a technique for true ACLs that support a list of users and groups with their permissions. Unfortunately, this is not widely supported nor supported exactly the same way across Unix-like systems. Stock Linux 2.2, for example, has neither ACLs nor POSIX capability values in the filesystem.

It's worth noting that in Linux, the Linux ext2 filesystem by default reserves a small amount of space for the root user. This is a partial defense against denial-of-service attacks; even if a user fills a disk that is shared with the root user, the root user has a little space left over (e.g., for critical functions). The default is 5% of the filesystem space; see `mke2fs(8)`, in particular its “-m” option.

### 3.2.2. Creation Time Initial Values

At creation time, the following rules apply. On most Unix systems, when a new filesystem object is created via `creat(2)` or `open(2)`, the FSO UID is set to the process' EUID and the FSO's GID is set to the process' EGID. Linux works slightly differently due to its FSUID extensions; the FSO's UID is set to the process' FSUID, and the FSO GID is set to the process' FSGUID; if the containing directory's `setgid` bit is set or the filesystem's “GRPID” flag is set, the FSO GID is actually set to the GID of the containing directory. Many systems, including Sun Solaris and Linux, also support the `setgid` directory extensions. As noted earlier, this special case supports “project” directories: to make a “project” directory, create a special group for the project, create a

directory for the project owned by that group, then make the directory setgid: files placed there are automatically owned by the project. Similarly, if a new subdirectory is created inside a directory with the setgid bit set (and the filesystem GRPID isn't set), the new subdirectory will also have its setgid bit set (so that project subdirectories will “do the right thing”); in all other cases the setgid is clear for a new file. This is the rationale for Red Hat Linux's “user-private group” scheme, in which every user is a member of a “private” group with just them as members, so their defaults can permit the group to read and write any file (since they're the only member of the group). Thus, when the file's group membership is transferred this way, read and write privileges are transferred too. FSO basic access control values (read, write, execute) are computed from (requested values & ~ umask of process). New files always start with a clear sticky bit and clear setuid bit.

### 3.2.3. Changing Access Control Attributes

You can set most of these values with `chmod(2)`, `fchmod(2)`, or `chmod(1)` but see also `chown(1)`, and `chgrp(1)`. In Linux, some the Linux-unique attributes are manipulated using `chattr(1)`.

Note that in Linux, only root can change the owner of a given file. Some Unix-like systems allow ordinary users to transfer ownership of their files to another, but this causes complications and is forbidden by Linux. For example, if you're trying to limit disk usage, allowing such operations would allow users to claim that large files actually belonged to some other “victim”.

### 3.2.4. Using Access Control Attributes

Under Linux and most Unix-like systems, reading and writing attribute values are only checked when the file is opened; they are not re-checked on every read or write. Still, a large number of calls do check these attributes, since the filesystem is so central to Unix-like systems. Calls that check these attributes include `open(2)`, `creat(2)`, `link(2)`, `unlink(2)`, `rename(2)`, `mknod(2)`, `symlink(2)`, and `socket(2)`.

### 3.2.5. Filesystem Hierarchy

Over the years conventions have been built on “what files to place where”. Where possible, please follow conventional use when placing information in the hierarchy. For example, place global configuration information in /etc. The Filesystem Hierarchy Standard (FHS) tries to define these conventions in a logical manner, and is widely used by Linux systems. The FHS is an update to the previous Linux Filesystem Structure standard (FSSTND), incorporating lessons learned and approaches from Linux, BSD, and System V systems. See <http://www.pathname.com/fhs> for more information about the FHS. A summary of these conventions is in hier(5) for Linux and hier(7) for Solaris. Sometimes different conventions disagree; where possible, make these situations configurable at compile or installation time.

## 3.3. System V IPC

Many Unix-like systems, including Linux and System V systems, support System V interprocess communication (IPC) objects. Indeed System V IPC is required by the Open Group’s Single UNIX Specification, Version 2 [Open Group 1997]. System V IPC objects can be one of three kinds: System V message queues, semaphore sets, and shared memory segments. Each such object has the following attributes:

- read and write permissions for each of creator, creator group, and others.
- creator UID and GID - UID and GID of the creator of the object.
- owning UID and GID - UID and GID of the owner of the object (initially equal to the creator UID).

When accessing such objects, the rules are as follows:

- if the process has root privileges, the access is granted.

- if the process' EUID is the owner or creator UID of the object, then the appropriate creator permission bit is checked to see if access is granted.
- if the process' EGID is the owner or creator GID of the object, or one of the process' groups is the owning or creating GID of the object, then the appropriate creator group permission bit is checked for access.
- otherwise, the appropriate "other" permission bit is checked for access.

Note that root, or a process with the EUID of either the owner or creator, can set the owning UID and owning GID and/or remove the object. More information is available in `ipc(5)`.

## 3.4. Sockets and Network Connections

Sockets are used for communication, particularly over a network. Sockets were originally developed by the BSD branch of Unix systems, but they are generally portable to other Unix-like systems: Linux and System V variants support sockets as well, and socket support is required by the Open Group's Single Unix Specification [Open Group 1997]. System V systems traditionally used a different (incompatible) network communication interface, but it's worth noting that systems like Solaris include support for sockets. `Socket(2)` creates an endpoint for communication and returns a descriptor, in a manner similar to `open(2)` for files. The parameters for socket specify the protocol family and type, such as the Internet domain (TCP/IP version 4), Novell's IPX, or the "Unix domain". A server then typically calls `bind(2)`, `listen(2)`, and `accept(2)` or `select(2)`. A client typically calls `bind(2)` (though that may be omitted) and `connect(2)`. See these routine's respective man pages for more information. It can be difficult to understand how to use sockets from their man pages; you might want to consult other papers such as Hall "Beej" [1999] to learn how these calls are used together.

The "Unix domain sockets" don't actually represent a network protocol; they can only connect to sockets on the same machine. (at the time of this writing for the standard Linux kernel). When used as a stream, they are fairly similar to named pipes, but with

significant advantages. In particular, Unix domain socket is connection-oriented; each new connection to the socket results in a new communication channel, a very different situation than with named pipes. Because of this property, Unix domain sockets are often used instead of named pipes to implement IPC for many important services. Just like you can have unnamed pipes, you can have unnamed Unix domain sockets using `socketpair(2)`; unnamed Unix domain sockets are useful for IPC in a way similar to unnamed pipes.

There are several interesting security implications of Unix domain sockets. First, although Unix domain sockets can appear in the filesystem and can have `stat(2)` applied to them, you can't use `open(2)` to open them (you have to use the `socket(2)` and friends interface). Second, Unix domain sockets can be used to pass file descriptors between processes (not just the file's contents). This odd capability, not available in any other IPC mechanism, has been used to hack all sorts of schemes (the descriptors can basically be used as a limited version of the "capability" in the computer science sense of the term). File descriptors are sent using `sendmsg(2)`, where the `msg` (message)'s field `msg_control` points to an array of control message headers (field `msg_controllen` must specify the number of bytes contained in the array). Each control message is a struct `cmsghdr` followed by data, and for this purpose you want the `cmsg_type` set to `SCM_RIGHTS`. A file descriptor is retrieved through `recvmsg(2)` and then tracked down in the analogous way. Frankly, this feature is quite baroque, but it's worth knowing about.

Linux 2.2 supports an addition feature in Unix domain sockets: you can acquire the peer's "credentials" (the `pid`, `uid`, and `gid`). Here's some sample code:

```
/* fd= file descriptor of Unix domain socket connected
   to the client you wish to identify */

struct ucred cr;
int cl=sizeof(cr);

if (getsockopt(fd, SOL_SOCKET, SO_PEERCRED, &cr, &cl)==0) {
    printf("Peer's pid=%d, uid=%d, gid=%d\n",
          cr.pid, cr.uid, cr.gid);
}
```

Standard Unix convention is that binding to TCP and UDP local port numbers less than 1024 requires root privilege, while any process can bind to an unbound port number of 1024 or greater. Linux follows this convention, more specifically, Linux requires a process to have the capability `CAP_NET_BIND_SERVICE` to bind to a port number less than 1024; this capability is normally only held by processes with an `euid` of 0. The adventurous can check this in Linux by examining its Linux's source; in Linux 2.2.12, it's file `/usr/src/linux/net/ipv4/af_inet.c`, function `inet_bind()`.

## 3.5. Signals

Signals are a simple form of “interruption” in the Unix-like OS world, and are an ancient part of Unix. A process can set a “signal” on another process (say using `kill(1)` or `kill(2)`), and that other process would receive and handle the signal asynchronously. For a process to have permission to send a signal to some other process, the sending process must either have root privileges, or the real or effective user ID of the sending process must equal the real or saved set-user-ID of the receiving process.

Although signals are an ancient part of Unix, they've had different semantics in different implementations. Basically, they involve questions such as “what happens when a signal occurs while handling another signal”? The older Linux libc 5 used a different set of semantics for some signal operations than the newer GNU libc libraries. For more information, see the glibc FAQ (on some systems a local copy is available at `/usr/doc/glibc-*/FAQ`).

For new programs, just use the POSIX signal system (which in turn was based on BSD work); this set is widely supported and doesn't have the problems that some of the older signal systems did. The POSIX signal system is based on using the `sigset_t` datatype, which can be manipulated through a set of operations: `sigemptyset()`, `sigfillset()`, `sigaddset()`, `sigdelset()`, and `sigismember()`. You can read about these in `sigsetops(3)`. Then use `sigaction(2)`, `sigaction(2)`, `sigprocmask(2)`, `sigpending(2)`, and `sigsuspend(2)` to set up an manipulate signal handling (see their man pages for more information).

In general, make any signal handlers very short and simple, and look carefully for race conditions. Signals, since they are by nature asynchronous, can easily cause race

conditions.

A common convention exists for servers: if you receive SIGHUP, you should close any log files, reopen and reread configuration files, and then re-open the log files. This supports reconfiguration without halting the server and log rotation without data loss. If you are writing a server where this convention makes sense, please support it.

## 3.6. Quotas and Limits

Many Unix-like systems have mechanisms to support filesystem quotas and process resource limits. This certainly includes Linux. These mechanisms are particularly useful for preventing denial of service attacks; by limiting the resources available to each user, you can make it hard for a single user to use up all the system resources. Be careful with terminology here, because both filesystem quotas and process resource limits have “hard” and “soft” limits but the terms mean slightly different things.

You can define storage (filesystem) quota limits on each mountpoint for the number of blocks of storage and/or the number of unique files (inodes) that can be used, and you can set such limits for a given user or a given group. A “hard” quota limit is a never-to-exceed limit, while a “soft” quota can be temporarily exceeded. See `quota(1)`, `quotactl(2)`, and `quotaon(8)`.

The `rlimit` mechanism supports a large number of process quotas, such as file size, number of child processes, number of open files, and so on. There is a “soft” limit (also called the current limit) and a “hard limit” (also called the upper limit). The soft limit cannot be exceeded at any time, but through calls it can be raised up to the value of the hard limit. See `getrlimit()`, `setrlimit()`, and `getrusage()`. Note that there are several ways to set these limits, including the PAM module `pam_limits`.

## 3.7. Dynamically Linked Libraries

Practically all programs depend on libraries to execute. In most modern Unix-like

systems, including Linux, programs are by default compiled to use *dynamically linked libraries* (DLLs). That way, you can update a library and all the programs using that library will use the new (hopefully improved) version if they can.

Dynamically linked libraries are typically placed in one a few special directories. The usual directories include `/lib`, `/usr/lib`, `/lib/security` for PAM modules, `/usr/X11R6/lib` for X-windows, and `/usr/local/lib`.

There are special conventions for naming libraries and having symbolic links for them, with the result that you can update libraries and still support programs that want to use old, non-backward-compatible versions of those libraries. There are also ways to override specific libraries or even just specific functions in a library when executing a particular program. This is a real advantage of Unix-like systems over Windows-like systems; I believe Unix-like systems have a much better system for handling library updates, one reason that Unix and Linux systems are reputed to be more stable than Windows-based systems.

On GNU glibc-based systems, including all Linux systems, the list of directories automatically searched during program start-up is stored in the file `/etc/ld.so.conf`. Many Red Hat-derived distributions don't normally include `/usr/local/lib` in the file `/etc/ld.so.conf`. I consider this a bug, and adding `/usr/local/lib` to `/etc/ld.so.conf` is a common "fix" required to run many programs on Red Hat-derived systems. If you want to just override a few functions in a library, but keep the rest of the library, you can enter the names of overriding libraries (`.o` files) in `/etc/ld.so.preload`; these "preloading" libraries will take precedence over the standard set. This preloading file is typically used for emergency patches; a distribution usually won't include such a file when delivered. Searching all of these directories at program start-up would be too time-consuming, so a caching arrangement is actually used. The program `ldconfig(8)` by default reads in the file `/etc/ld.so.conf`, sets up the appropriate symbolic links in the dynamic link directories (so they'll follow the standard conventions), and then writes a cache to `/etc/ld.so.cache` that's then used by other programs. So, `ldconfig` has to be run whenever a DLL is added, when a DLL is removed, or when the set of DLL directories changes; running `ldconfig` is often one of the steps performed by package managers when installing a library. On start-up, then, a program uses the dynamic loader to read the file `/etc/ld.so.cache` and then load the libraries it needs.



Various environment variables can control this process, and in fact there are environment variables that permit you to override this process (so, for example, you can temporarily substitute a different library for this particular execution). In Linux, the environment variable `LD_LIBRARY_PATH` is a colon-separated set of directories where libraries should be searched for first, before the standard set of directories; this is useful when debugging a new library or using a nonstandard library for special purposes. The variable `LD_PRELOAD` lists object files with functions that override the standard set, just as `/etc/ld.so.preload` does.

Permitting user control over dynamically linked libraries would be disastrous for `setuid/setgid` programs if special measures weren't taken. Therefore, in the GNU glibc implementation, if the program is `setuid` or `setgid` these variables (and other similar variables) are ignored or greatly limited in what they can do. The GNU glibc library determines if a program is `setuid` or `setgid` by checking the program's credentials; if the `uid` and `euid` differ, or the `gid` and the `egid` differ, the library presumes the program is `setuid/setgid` (or descended from one) and therefore greatly limits its abilities to control linking. If you load the GNU glibc libraries, you can see this; see especially the files `elf/rtld.c` and `sysdeps/generic/dl-sysdep.c`. This means that if you cause the `uid` and `gid` to equal the `euid` and `egid`, and then call a program, these variables will have full effect. Other Unix-like systems handle the situation differently but for the same reason: a `setuid/setgid` program should not be unduly affected by the environment variables set.

## 3.8. Audit

Different Unix-like systems handle auditing differently. In Linux, the most common "audit" mechanism is `syslogd(8)`, usually working in conjunction with `klogd(8)`. You might also want to look at `wtmp(5)`, `utmp(5)`, `lastlog(8)`, and `acct(2)`. Some server programs (such as the Apache web server) also have their own audit trail mechanisms. According to the FHS, audit logs should be stored in `/var/log` or its subdirectories.

## **3.9. PAM**

Sun Solaris and nearly all Linux systems use the Pluggable Authentication Modules (PAM) system for authentication. PAM permits run-time configuration of authentication methods (e.g., use of passwords, smart cards, etc.). PAM will be discussed more fully later in this document.

# Chapter 4. Validate All Input

*Wisdom will save you from the  
ways of wicked men, from men  
whose words are perverse...  
Proverbs 2:12 (NIV)*

Some inputs are from untrustable users, so those inputs must be validated (filtered) before being used. You should determine what is legal and reject anything that does not match that definition. Do not do the reverse (identify what is illegal and reject those cases), because you are likely to forget to handle an important case. Limit the maximum character length (and minimum length if appropriate), and be sure to not lose control when such lengths are exceeded (see Chapter 5 for more about this kind of problem, called a buffer overflow).

For strings, identify the legal characters or legal patterns (e.g., as a regular expression) and reject anything not matching that form. There are special problems when strings contain control characters (especially linefeed or NIL) or shell metacharacters; it is often best to “escape” such metacharacters immediately when the input is received so that such characters are not accidentally sent. CERT goes further and recommends escaping all characters that aren’t in a list of characters not needing escaping [CERT 1998, CMU 1998]. See Section 7.1 for more information on limiting call-outs.

Limit all numbers to the minimum (often zero) and maximum allowed values. Filenames should be checked; usually you will want to not include “.” (higher directory) as a legal value. In filenames it’s best to prohibit any change in directory, e.g., by not including “/” in the set of legal characters. A full email address checker is actually quite complicated, because there are legacy formats that greatly complicate validation if you need to support all of them; see `mailaddr(7)` and IETF RFC 822 [RFC 822] for more information if such checking is necessary.

The legal character patterns must not include characters or character sequences that have special meaning to the program internals or the eventual output unless you account for them. In particular, if you store data (internally or externally) in delimited strings, make sure that the delimiters are not permitted data values. Here are two

common cases:

- A character sequence may have special meaning to the program's internal storage format. A number of programs store data in comma (,) or colon (:) delimited text files; inserting such values in the input can be problem unless the program accounts for it. Other characters often causing these problems include single and double quotes (used for surrounding strings) and the less-than sign (used in SGML, XML, and HTML to indicate a tag's beginning). Most data formats have an escape sequence to handle these cases; use it, or filter such data on input.
- A character sequence may have special meaning if sent back out to the user. Another common case is permitting HTML tags in data input that will later be posted to other readers (e.g., in a guestbook or "reader comment" area). These tags can be used by malicious users to attack other users by inserting Java references (including references to hostile applets), DHTML tags, early document endings (via `</HTML>`), absurd font size requests, and so on, causing anything from unreadable pages to destructive attacks. It's safest to strip or escape all HTML tags, but at least identify a list of "safe" HTML commands and only permit those commands. Common safe HTML tags that might be useful for guestbook or other applications supporting short comments include `<P>` (paragraph), `<B>` (bold), `<I>` (italics), `<EM>` (emphasis), `<STRONG>` (strong emphasis), `<PRE>` (preformatted text), `<BR>` (forced line break), and `<A HREF="safe URI">` (hypertext link), as well as all their ending tags. You might even consider supporting the list-oriented tags, such as `<OL>` (ordered list), `<UL>` (unordered list), and `<LI>` (list item). It's tricky to define "safe URI"; I'd suggest a pattern like `"(http|ftp):/[A-Za-z0-9._]+"` (this allows `..`, which is often fine in this application, but note that it intentionally prevents most query formats and other schemes like `mailto`). There are more HTML tags, but after a certain point you're really permitting full publishing (in which case you need to trust them or perform more serious checking than will be described here). You really should check if the HTML commands are properly nested (though supporting an implied `</P>` where not provided before a `<P>` would be fine), and if you support list tags further checking is warranted.

These tests should usually be centralized in one place so that the validity tests can be

easily examined for correctness later.

Make sure that your validity test is actually correct; this is particularly a problem when checking input that will be used by another program (such as a filename, email address, or URL). Often these tests have subtle errors, producing the so-called “deputy problem” (where the checking program makes different assumptions than the program that actually uses the data).

While parsing user input, it’s a good idea to temporarily drop all privileges, or even create separate processes (with the parser having permanently dropped privileges, and the other process performing security checks against the parser requests). This is especially true if the parsing task is complex (e.g., if you use a lex-like or yacc-like tool), or if the programming language doesn’t protect against buffer overflows (e.g., C and C++). See Section 6.2 for more information on minimizing privileges.

The following subsections discuss different kinds of inputs to a program; note that input includes process state such as environment variables, umask values, and so on. Not all inputs are under the control of an untrusted user, so you need only worry about those inputs that are.

## 4.1. Command line

Many programs use the command line as an input interface, accepting input by being passed arguments. A `setuid/setgid` program has a command line interface provided to it by an untrusted user, so it must defend itself. Users have great control over the command line (through calls such as the `execve(3)` call). Therefore, `setuid/setgid` programs must validate the command line inputs and must not trust the name of the program reported by command line argument zero (the user can set it to any value including `NULL`).

## 4.2. Environment Variables

By default, environment variables are inherited from a process’ parent. However, when

a program executes another program, the calling program can set the environment variables to arbitrary values. This is dangerous to `setuid/setgid` programs, because their invoker can completely control the environment variables they're given. Since they are usually inherited, this also applies transitively; a secure program might call some other program and, without special measures, would pass potentially dangerous environment variables values on to the program it calls.

### 4.2.1. Some Environment Variables are Dangerous

Some environment variables are dangerous because many libraries and programs are controlled by environment variables in ways that are obscure, subtle, or undocumented. For example, the `IFS` variable is used by the `sh` and `bash` shell to determine which characters separate command line arguments. Since the shell is invoked by several low-level calls (like `system(3)` and `popen(3)` in C, or the back-tick operator in Perl), setting `IFS` to unusual values can subvert apparently-safe calls. This behavior is documented in `bash` and `sh`, but it's obscure; many long-time users only know about `IFS` because of its use in breaking security, not because it's actually used very often for its intended purpose. What is worse is that not all environment variables are documented, and even if they are, those other programs may change and add dangerous environment variables. Thus, the only real solution (described below) is to select the ones you need and throw away the rest.

### 4.2.2. Environment Variable Storage Format is Dangerous

Normally, programs should use the standard access routines to access environment variables. For example, in C, you should get values using `getenv(3)`, set them using the POSIX standard routine `putenv(3)` or the BSD extension `setenv(3)` and eliminate environment variables using `unsetenv(3)`. I should note here that `setenv(3)` is implemented in Linux, too. However, crackers need not be so nice; crackers can directly control the environment variable data area passed to a program using `execve(2)`. This permits some nasty attacks, which can only be understood by understanding how environment variables really work. In Linux, you can see `environ(5)` for a summary

how about environment variables really work. In short, environment variables are internally stored as a pointer to an array of pointers to characters; this array is stored in order and terminated by a NULL pointer (so you'll know when the array ends). The pointers to characters, in turn, each point to a NIL-terminated string value of the form "NAME=value". This has several implications, for example, environment variable names can't include the equal sign, and neither the name nor value can have embedded NIL characters. However, a more dangerous implication of this format is that it allows multiple entries with the same variable name, but with different values (e.g., more than one value for SHELL). While typical command shells prohibit doing this, a locally-executing cracker can create such a situation using `execve(2)`.

The problem with this storage format (and the way it's set) is that a program might check one of these values (to see if it's valid) but actually use a different one. In Linux, the GNU glibc libraries try to shield programs from this; glibc 2.1's implementation of `getenv` will always get the first matching entry, `setenv` and `putenv` will always set the first matching entry, and `unsetenv` will actually unset *all* of the matching entries (congratulations to the GNU glibc implementors for implementing `unsetenv` this way!). However, some programs go directly to the `environ` variable and iterate across all environment variables; in this case, they might use the last matching entry instead of the first one. As a result, if checks were made against the first matching entry instead, but the actual value used is the last matching entry, a cracker can use this fact to circumvent the protection routines.

### 4.2.3. The Solution - Extract and Erase

For secure `setuid/setgid` programs, the short list of environment variables needed as input (if any) should be carefully extracted. Then the entire environment should be erased, followed by resetting a small set of necessary environment variables to safe values. There really isn't a better way if you make any calls to subordinate programs; there's no practical method of listing "all the dangerous values". Even if you reviewed the source code of every program you call directly or indirectly, someone may add new undocumented environment variables after you write your code, and one of them may be exploitable.

The simple way to erase the environment is by setting the global variable *environ* to NULL. The global variable *environ* is defined in `<unistd.h>`; C/C++ users will want to `#include` this header file. You will need to manipulate this value before spawning threads, but that's rarely a problem, since you want to do these manipulations very early in the program's execution (usually before threads are spawned). Another way is to use the undocumented `clearenv()` function. `clearenv()` has an odd history; it was supposed to be defined in POSIX.1, but somehow never made it into that standard. However, `clearenv()` is defined in POSIX.9 (the Fortran 77 bindings to POSIX), so there is a quasi-official status for it. In Linux, `clearenv()` is defined in `<stdlib.h>`, but before using `#include` to include it you must make sure that `__USE_MISC` is `#defined`. A somewhat more "official" way to cause `__USE_MISC` to be defined is to `#define` either `_SVID_SOURCE` or `_BSD_SOURCE`, then `#include <features.h>` - these are the official feature test macros. To be honest, I use the approach of setting "environ"; manipulating such low-level components is possibly non-portable, but it assures you that you get a clean (and safe) environment. In the rare case where you need later access to the entire set of variables, you could save the "environ" variable's value somewhere, but this is rarely necessary; nearly all programs need only a few values, and the rest can be dropped.

One value you'll almost certainly re-add is `PATH`, the list of directories to search for programs; `PATH` should *not* include the current directory and usually be something simple like `"/bin:/usr/bin"`. Typically you'll also set `IFS` (to its default of `"\t\n"`) and `TZ` (timezone). Linux won't die if you don't supply either `IFS` or `TZ`, but some System V based systems have problems if you don't supply a `TZ` value, and it's rumored that some shells need the `IFS` value set. In Linux, see `environ(5)` for a list of common environment variables that you *might* want to set.

If you really need user-supplied values, check the values first (to ensure that the values match a pattern for legal values and that they are within some reasonable maximum length). Ideally there would be some standard trusted file in `/etc` with the information for "standard safe environment variable values", but at this time there's no standard file defined for this purpose. For something similar, you might want to examine the PAM module `pam_env` on those systems which have that module.

If you're programming a `setuid/setgid` program in a language that doesn't allow you to reset the environment directly, one approach is to create a "wrapper" program. The



wrapper sets the environment program to safe values, and then calls the other program. Beware: make sure the wrapper will actually invoke the intended program; if it's an interpreted program, make sure there's no race condition possible that would allow the interpreter to load a different program than the one that was granted the special `setuid/setgid` privileges.

## 4.3. File Descriptors

A program is passed a set of “open file descriptors”, that is, pre-opened files. A `setuid/setgid` program must deal with the fact that the user gets to select what files are open and to what (within their permission limits). A `setuid/setgid` program must not assume that opening a new file will always open into a fixed file descriptor id. It must also not assume that standard input (`stdin`), standard output (`stdout`), and standard error (`stderr`) refer to a terminal or are even open.

The rationale behind this is easy; since an attacker can open or close a file descriptor before starting the program, the attacker could create an unexpected situation. If the attacker closes the standard output, when the program opens the next file it will be opened as though it were standard output, and then it will send all standard output to that file as well. Some C libraries will automatically open `stdin`, `stdout`, and `stderr` if they aren't already open (to `/dev/null`), but this isn't true on all Unix-like systems.

## 4.4. File Contents

If a program takes directions from a file, it must not trust that file specially unless only a trusted user can control its contents. Usually this means that an untrusted user must not be able to modify the file, its directory, or any of its ancestor directories. Otherwise, the file must be treated as suspect.

If the directions in the file are supposed to be from an untrusted user, then make sure that the inputs from the file are protected as describe throughout this document. In

particular, check that values match the set of legal values, and that buffers are not overflowed.

## 4.5. Web-Based Applications (Especially CGI Scripts)

Web-based applications (such as CGI scripts) run on some trusted server and must get their input data somehow through the web. Since the input data generally come from untrusted users, this input data must be validated. For example, CGI scripts are passed this information through a standard set of environment variables and through standard input. The rest of this text will specifically discuss CGI, because it's the most common technique for implementing dynamic web content, but the general issues are the same for most other dynamic web content techniques.

One additional complication is that many CGI inputs are provided in so-called "URL-encoded" format, that is, some values are written in the format %HH where HH is the hexadecimal code for that byte. You or your CGI library must handle these inputs correctly by URL-decoding the input and then checking if the resulting byte value is acceptable. You must correctly handle all values, including problematic values such as %00 (NIL) and %0A (newline). Don't decode inputs more than once, or input such as "%2500" will be mishandled (the %25 would be translated to "%", and the resulting "%00" would be erroneously translated to the NIL character).

CGI scripts are commonly attacked by including special characters in their inputs; see the comments above.

Some HTML forms include client-side checking to prevent some illegal values. This checking can be helpful for the user but is useless for security, because attackers can send such "illegal" values directly to the web server. In general, servers must perform all their own input checking, because they cannot trust clients to do this securely (clients are generally not "trustworthy channels" for such information). See Section 6.7 for more information on trustworthy channels.

## 4.6. Other Inputs

Programs must ensure that all inputs are controlled; this is particularly difficult for `setuid/setgid` programs because they have so many such inputs. Other inputs programs must consider include the current directory, signals, memory maps (`mmaps`), System V IPC, and the `umask` (which determines the default permissions of newly-created files). Consider explicitly changing directories (using `chdir(2)`) to an appropriately fully named directory at program startup.

## 4.7. Human Language (Locale) Selection

As more people have computers and the Internet available to them, there has been increasing pressure for programs to support multiple human languages and cultures. This combination of language and other cultural factors is usually called a “locale”. The process of modifying a program so it can support multiple locales is called “internationalization” (i18n), and the process of providing the information for a particular locale to a program is called “localization” (l10n).

Overall, internationalization is a good thing, but this process provides another opportunity for a security exploit. Since a potentially untrusted user provides information on the desired locale, locale selection becomes another input that, if not properly protected, can be exploited.

### 4.7.1. How Locales are Selected

In locally-run programs (including `setuid/setgid` programs), locale information is provided by an environment variable. Thus, like all other environment variables, these values must be extracted and checked against valid patterns before use.

For web applications, this information can be obtained from the web browser (via the `Accept-Language` request header). However, since not all web browsers properly pass this information (and not all users configure their browsers properly), this is used less often than you might think. Often, the language requested in a web browser is simply

passed in as a form value. Again, these values must be checked for validity before use, as with any other form value.

In either case, locale information is really just a special case of input discussed in the previous sections. However, because this input is so rarely considered, I'm discussing it separately. In particular, when combined with format strings (discussed later), user-controlled strings can permit attackers to force other programs to run arbitrary instructions, corrupt data, and do other unfortunate actions.

## 4.7.2. Locale Support Mechanisms

There are two major library interfaces for supporting locale-selected messages on Unix-like systems, one called “catgets” and the other called “gettext”. In the catgets approach, every string is assigned a unique number, which is used as an index into a table of messages. In contrast, in the gettext approach, a string (usually in English) is used to look up a table that translates the original string. catgets(3) is an accepted standard (via the X/Open Portability Guide, Volume 3 and Single Unix Specification), so it's possible your program uses it. The “gettext” interface is not an official standard, (though it was originally a UniForum proposal), but I believe it's the more widely used interface (it's used by Sun and essentially all GNU programs).

In theory, catgets should be slightly faster, but this is at best marginal on today's machines, and the bookkeeping effort to keep unique identifiers valid in catgets() makes the gettext() interface much easier to use. I'd suggest using gettext(), just because it's easier to use. However, don't take my word for it; see GNU's documentation on gettext (info:gettext#catgets) for a longer and more descriptive comparison.

The catgets(3) call (and its associated catopen(3) call) in particular is vulnerable to security problems, because the environment variable NLSPATH can be used to control the filenames used to acquire internationalized messages. The GNU C library ignores NLSPATH for setuid/setgid programs, which helps, but that doesn't protect programs running on other implementations, nor other programs (like CGI scripts) which don't “appear” to require such protection.

The widely-used “gettext” interface is at least not vulnerable to a malicious NLSPATH

setting to my knowledge. However, it appears likely to me that malicious settings of LC\_ALL or LC\_MESSAGES could cause problems. Also, if you use gettext's bindtextdomain() routine in its file cat-compat.c, that does depend on NLSPATH.

### 4.7.3. Legal Values

For the moment, if you must permit untrusted users to set information on their desired locales, make sure the provided internationalization information meets a narrow filter that only permits legitimate locale names. For user programs (especially setuid/setgid programs), these values will come in via NLSPATH, LANGUAGE, LANG, the old LINGUAS, LC\_ALL, and the other LC\_\* values (especially LC\_MESSAGES, but also including LC\_COLLATE, LC\_CTYPE, LC\_MONETARY, LC\_NUMERIC, and LC\_TIME). For web applications, this user-requested set of language information would be done via the Accept-Language request header or a form value (the application should indicate the actual language setting of the data being returned via the Content-Language heading). You can check this value as part of your environment variable filtering if your users can set your environment variables (i.e., setuid/setgid programs) or as part of your input filtering (e.g., for CGI scripts). The GNU C library "glibc" doesn't accept some values of LANG for setuid/setgid programs (in particular anything with "/"), but errors have been found in that filtering (e.g., Red Hat released an update to fix this error in glibc on September 1, 2000). This kind of filtering isn't required by any standard, so you're safer doing this filtering yourself. I have not found any guidance on filtering language settings, so here are my suggestions based on my own research into the issue.

First, a few words about the legal values of these settings. Language settings are generally set using the standard tags defined in IETF RFC 1766 (which uses two-letter country codes as its basic tag, followed by an optional subtag separated by a dash; I've found that environment variable settings use the underscore instead). However, some find this insufficiently flexible, so three-letter country codes may soon be used as well. Also, there are two major not-quite-compatible extended formats, the X/Open Format and the CEN Format (European Community Standard); you'd like to permit both. Typical values include "C" (the C locale), "EN" (English), and "FR\_fr" (French using

the territory of France's conventions). Also, so many people use nonstandard names that programs have had to develop "alias" systems to cope with them (for GNU gettext, see /usr/share/locale/locale.alias, and for X11, see /usr/lib/X11/locale/locale.alias); they should usually be permitted as well. Libraries like gettext() have to accept all these variants and find an appropriate value, where possible. One source of further information is FSF [1999]. However, a filter should not permit characters that aren't needed, in particular "/" (which might permit escaping out of the trusted directories) and "." (which might permit going up one directory). Other dangerous characters in NLSPATH include "%" (which indicates substitution) and ":" (which is the directory separator); the documentation I have for other machines suggests that some implementations may use them for other values, so it's safest to prohibit them.

#### 4.7.4. Bottom Line

In short, I suggest simply erasing or re-setting the NLSPATH, unless you have a trusted user supplying the value. For the Accept-Language heading in HTTP (if you use it), form values specifying the locale, and the environment variables LANGUAGE, LANG, the old LINGUAS, LC\_ALL, and the other LC\_\* values listed above, filter the locales from untrusted users to permit null (empty) values or to only permit values that match in total this regular expression:

```
[A-Za-z][A-Za-z0-9_+@-\.\ ]*
```

I haven't found any legitimate locale which doesn't match this pattern, but this pattern does appear to protect against locale attacks. Of course, there's no guarantee that there are messages available in the requested locale, but in such a case these routines will fall back to the default messages (usually in English), which at least is not a security problem.

Of course, languages cannot be supported without a standard way to represent their written symbols, which brings us to the issue of character encoding.

## 4.8. Character Encoding

For many years Americans have been using the ASCII encoding of characters, permitting easy exchange of English texts. Unfortunately, ASCII is completely inadequate in handling the character sets of most other languages. For many years different countries have adopted different techniques for exchanging text in different languages. More recently, ISO has developed ISO 10646, a single 31-bit encoding for all of the world's characters termed the Universal Character Set (UCS). Characters fitting into 16 bits (the first 65536 values of the UCS) are termed the "Basic Multilingual Plane" (BMP), and the BMP is intended to cover nearly all spoken languages. The Unicode forum develops the Unicode standard, which concentrates on the 16-bit set and adds some additional conventions to aid interoperability.

However, most software is not designed to handle 16 bit or 32 bit characters, so a special format called "UTF-8" was developed to encode these potentially international characters in a format more easily handled by existing programs and libraries. UTF-8 is defined, among other places, in IETF RFC 2279, so it's a well-defined standard that can be freely read and used. UTF-8 is a variable-width encoding; characters numbered 0 to 0x7f (127) encode to themselves as a single byte, while characters with larger values are encoded into 2 to 6 bytes of information (depending on their value). The encoding has been specially designed to have the following nice properties (this information is from the RFC and Linux utf-8 man page):

- The classical US ASCII characters (0 to 0x7f) encode as themselves, so files and strings which contain only 7-bit ASCII characters have the same encoding under both ASCII and UTF-8. This is fabulous for backwards compatibility with the many existing U.S. programs and data files.
- All UCS characters beyond 0x7f are encoded as a multibyte sequence consisting only of bytes in the range 0x80 to 0xfd. This means that no ASCII byte can appear as part of another character. Many other encodings permit characters such as an embedded NIL, causing programs to fail.
- It's easy to convert between UTF-8 and a 2-byte or 4-byte fixed-width representations of characters (these are called UCS-2 and UCS-4 respectively).

- The lexicographic sorting order of UCS-4 strings is preserved, and the Boyer-Moore fast search algorithm can be used directly with UTF-8 data.
- All possible  $2^{31}$  UCS codes can be encoded using UTF-8.
- The first byte of a multibyte sequence which represents a single non-ASCII UCS character is always in the range 0xc0 to 0xfd and indicates how long this multibyte sequence is. All further bytes in a multibyte sequence are in the range 0x80 to 0xbf. This allows easy resynchronization; if a byte is missing, it's easy to skip forward to the "next" character, and it's always easy to skip forward and back to the "next" or "preceding" character.

In short, the UTF-8 transformation format is becoming a dominant method for exchanging international text information because it can support all of the world's languages, yet it is backward compatible with U.S. ASCII files as well as having other nice properties. For many purposes I recommend its use, particularly when storing data in a "text" file.

The reason to mention UTF-8 is that some byte sequences are not legal UTF-8, and this might be an exploitable security hole. The RFC notes the following:

Implementors of UTF-8 need to consider the security aspects of how they handle illegal UTF-8 sequences. It is conceivable that in some circumstances an attacker would be able to exploit an incautious UTF-8 parser by sending it an octet sequence that is not permitted by the UTF-8 syntax.

A particularly subtle form of this attack could be carried out against a parser which performs security-critical validity checks against the UTF-8 encoded form of its input, but interprets certain illegal octet sequences as characters. For example, a parser might prohibit the NUL character when encoded as the single-octet sequence 00, but allow the illegal two-octet sequence C0 80 and interpret it as a NUL character. Another example might be a parser which prohibits the octet sequence 2F 2E 2E 2F ("/./."), yet permits the illegal octet sequence 2F C0 AE 2E 2F.

A longer discussion about this is available at Markus Kuhn's *UTF-8 and Unicode FAQ for Unix/Linux* at <http://www.cl.cam.ac.uk/~mgk25/unicode.html>.

The UTF-8 character set is one case where it's possible to enumerate all illegal values (and prove that you've enumerated them all). If you need to determine if you have a



legal UTF-8 sequence, you need to check for two things: (1) is the initial sequence legal, and (2) if it is, is the first byte followed by the required number of valid continuation characters? Performing the first check is easy; the following is provably the complete list of all illegal UTF-8 initial sequences:

**Table 4-1. Illegal UTF-8 initial sequences**

UTF-8 Sequence	Reason for Illegality
10xxxxxx	illegal as initial byte of character (80..BF)
1100000x	illegal, overlong (C0 80..BF)
11100000 100xxxxx	illegal, overlong (E0 80..9F)
11110000 1000xxxx	illegal, overlong (F0 80..8F)
11111000 10000xxx	illegal, overlong (F8 80..87)
11111100 100000xx	illegal, overlong (FC 80..83)
1111111x	illegal; prohibited by spec

I should note that in some cases, you might want to cut slack (or use internally) the hexadecimal sequence C0 80. This is an overlong sequence that could represent ASCII NUL (NIL). Since C/C++ have trouble including a NIL character in an ordinary string, some people have taken to using this sequence when they want to represent NIL as part of the data stream; Java even enshrines the practice. Feel free to use C0 80 internally while processing data, but technically you really should translate this back to 00 before saving the data. Depending on your needs, you might decide to be “sloppy” and accept C0 80 as input in a UTF-8 data stream.

The second step is to check if the correct number of continuation characters are included in the string. If the first byte has the top 2 bits set, you count the number of “one” bits set after the top one, and then check that there are that many continuation bytes which begin with the bits “10”. So, binary 11100001 requires two more continuation bytes.

A related issue is that some phrases can be expressed in more than one way in ISO 10646/Unicode. For example, some accented characters can be represented as a single

character (with the accent) and also as a set of characters (e.g., the base character plus a separate composing accent). These two forms may appear identical. There's also a zero-width space that could be inserted, with the result that apparently-similar items are considered different. Beware of situations where such hidden text could interfere with the program.

## **4.9. Limit Valid Input Time and Load Level**

Place timeouts and load level limits, especially on incoming network data. Otherwise, an attacker might be able to easily cause a denial of service by constantly requesting the service.

# Chapter 5. Avoid Buffer Overflow

*An enemy will overrun the land;  
he will pull down your  
strongholds and plunder your  
fortresses.*

*Amos 3:11 (NIV)*

An extremely common security flaw is the “buffer overflow”. Technically, a buffer overflow is a problem with the program’s internal implementation, but it’s such a common and serious problem that I’ve placed this information in its own chapter. To give you an idea of how important this subject is, at the CERT, 9 of 13 advisories in 1998 and at least half of the 1999 advisories involved buffer overflows. An informal survey on Bugtraq found that approximately 2/3 of the respondents felt that buffer overflows were the leading cause of security vulnerability (the remaining respondents identified “misconfiguration” as the leading cause) [Cowan 1999]. This is an old, well-known problem, yet it continues to resurface [McGraw 2000].

A buffer overflow occurs when you write a set of values (usually a string of characters) into a fixed length buffer and write at least one value outside that buffer’s boundaries (usually past its end). A buffer overflow can occur when reading input from the user into a buffer, but it can also occur during other kinds of processing in a program.

If a secure program permits a buffer overflow, the overflow can often be exploited by an adversary. If the buffer is a local C variable, the overflow can be used to force the function to run code of an attacker’s choosing. This specific variation is often called a “stack smashing” attack. A buffer in the heap isn’t much better; attackers may be able to use such overflows to control other variables in the program. More details can be found from Aleph1 [1996], Mudge [1995], or the Nathan P. Smith’s “Stack Smashing Security Vulnerabilities” website at <http://destroy.net/machines/security/>.

Most high-level programming languages are essentially immune to this problem, either because they automatically resize arrays (e.g., Perl), or because they normally detect and prevent buffer overflows (e.g., Ada95). However, the C language provides no protection against such problems, and C++ can be easily used in ways to cause this

problem too.

## 5.1. Dangers in C/C++

C users must avoid using dangerous functions that do not check bounds unless they've ensured that the bounds will never get exceeded. Functions to avoid in most cases (or ensure protection) include the functions `strcpy(3)`, `strcat(3)`, `sprintf(3)` (with cousin `vsprintf(3)`), and `gets(3)`. These should be replaced with functions such as `strncpy(3)`, `strncat(3)`, `snprintf(3)`, and `fgets(3)` respectively, but see the discussion below. The function `strlen(3)` should be avoided unless you can ensure that there will be a terminating NIL character to find. The `scanf()` family (`scanf(3)`, `fscanf(3)`, `sscanf(3)`, `vscanf(3)`, `vsscanf(3)`, and `vfscanf(3)`) is often dangerous to use; do not use it to send data to a string without controlling the maximum length (the format `%s` is a particularly common problem). Other dangerous functions that may permit buffer overruns (depending on their use) include `realpath(3)`, `getopt(3)`, `getpass(3)`, `streadd(3)`, `strcpy(3)`, and `strtrns(3)`. You must be careful with `getwd(3)`; the buffer sent to `getwd(3)` must be at least `PATH_MAX` bytes long.

Unfortunately, `snprintf()`'s variants have additional problems. Officially, `snprintf()` is not a standard C function in the ISO 1990 (ANSI 1989) standard, though `sprintf()` is, so not all systems include `snprintf()`. Even worse, some systems' `snprintf()` do not actually protect against buffer overflows; they just call `sprintf` directly. Old versions of Linux's `libc4` depended on a "libbsd" that did this horrible thing, and I'm told that some old HP systems did the same. Linux's current version of `snprintf` is known to work correctly, that is, it does actually respect the boundary requested. The return value of `snprintf()` varies as well; the Single Unix Specification (SUS) version 2 and the upcoming C99 standard differ on what is returned by `snprintf()`. Finally, it appears that at least some versions of `snprintf` don't guarantee that its string will end in NIL; if the string is too long, it won't include NIL at all. Note that the `glib` library (the basis of GTK, and not the same as the GNU C library `glibc`) has a `g_snprintf()`, which has a consistent return semantic, always NIL-terminates, and most importantly always respects the buffer length.

## 5.2. Library Solutions in C/C++

One solution in C/C++ is to use library functions that do not have buffer overflow problems. The first subsection describes the “standard C library” solution, which can work but has its disadvantages. The next subsection describes the general security issues of both fixed length and dynamically reallocated approaches to buffers. The following subsections describe various alternative libraries, such as `strncpy` and `libmib`.

### 5.2.1. Standard C Library Solution

The “standard” solution to prevent buffer overflow in C is to use the standard C library calls that defend against these problems. This approach depends heavily on the standard library functions `strncpy(3)` and `strncat(3)`. If you choose this approach, beware: these calls have somewhat surprising semantics and are hard to use correctly. The function `strncpy(3)` does not `NIL`-terminate the destination string if the source string length is at least equal to the destination’s, so be sure to set the last character of the destination string to `NIL` after calling `strncpy(3)`. If you’re going to reuse the same buffer many times, an efficient approach is to tell `strncpy()` that the buffer is one character shorter than it actually is and set the last character to `NIL` once before use. Both `strncpy(3)` and `strncat(3)` require that you pass the amount of space left available, a computation that is easy to get wrong (and getting it wrong could permit a buffer overflow attack). Neither provide a simple mechanism to determine if an overflow has occurred. Finally, `strncpy(3)` has a significant performance penalty compared to the `strcpy(3)` it supposedly replaces, because *strncpy(3) `NIL`-fills the remainder of the destination*. I’ve gotten emails expressing surprise over this last point, but this is clearly stated in Kernighan and Ritchie second edition [Kernighan 1988, page 249], and this behavior is clearly documented in the man pages for Linux, FreeBSD, and Solaris. This means that just changing from `strcpy` to `strncpy` can cause a severe reduction in performance, for no good reason in most cases.

You can also use `sprintf()` while preventing buffer overflows, but you need to be careful when doing so; it’s so easy to misapply that it’s hard to recommend. The `sprintf` control string can contain various conversion specifiers (e.g., “%s”), and the control specifiers can have optional field width (e.g., “%10s”) and precision (e.g., “%.10s”)

specifications. These look quite similar (the only difference is a period) but they are very different. The field width only specifies a *minimum* length and is completely worthless for preventing buffer overflows. In contrast, the precision specification specifies the maximum length that that particular string may have in its output when used as a string conversion specifier - and thus it can be used to protect against buffer overflows. Note that the precision specification only specifies the total maximum length when dealing with a string; it has a different meaning for other conversion operations. If the size is given as "\*", then you can pass the maximum size as a parameter (e.g., the result of a sizeof() operation). This is most easily shown by an example - here's the wrong and right way to use sprintf() to protect against buffer overflows:

```
char buf[BUFFER_SIZE];
sprintf(buf, "%*s", sizeof(buf)-1, "long-string"); /* WRONG */
sprintf(buf, "%. *s", sizeof(buf)-1, "long-string"); /* RIGHT */
```

In theory, sprintf() should be very helpful because you can use it to specify complex formats. Sadly, it's easy to get things wrong with sprintf(). If the format is complex, you need to make sure that the destination is large enough for the largest possible size of the *entire* format, but the precision field only controls the size of one parameter. The "largest possible" value is often hard to determine when a complicated output is being created. If a program doesn't allocate quite enough space for the longest possible combination, a buffer overflow vulnerability may open up. Also, sprintf() appends a NUL to the destination after the entire operation is complete - this extra character is easy to forget and creates an opportunity for off-by-one errors. So, while this works, it can be painful to use in some circumstances.

## 5.2.2. Static and Dynamically Allocated Buffers

strncpy and friends are an example of statically allocated buffers, that is, once the buffer is allocated it stays a fixed size. The alternative is to dynamically reallocate buffer sizes as you need them. It turns out that both approaches have security implications.

There is a general security problem when using fixed-length buffers: the fact that the buffer is a fixed length may be exploitable. This is a problem with strncpy(3) and

`strncat(3)`, `snprintf(3)`, `strncpy(3)`, `strlcat(3)`, and other such functions. The basic idea is that the attacker sets up a really long string so that, when the string is truncated, the final result will be what the attacker wanted (instead of what the developer intended). Perhaps the string is concatenated from several smaller pieces; the attacker might make the first piece as long as the entire buffer, so all later attempts to concatenate strings do nothing. Here are some specific examples:

- Imagine code that calls `gethostbyname(3)` and, if successful, immediately copies `hostent->h_name` to a fixed-size buffer using `strncpy` or `snprintf`. Using `strncpy` or `snprintf` protects against an overflow of an excessively long fully-qualified domain name (FQDN), so you might think you're done. However, this could result in chopping off the end of the FQDN. This may be very undesirable, depending on what happens next.
- Imagine code that uses `strncpy`, `strncat`, `snprintf`, etc., to copy the full path of a filesystem object to some buffer. Further imagine that the original value was provided by an untrusted user, and that the copying is part of a process to pass a resulting computation to a function. Sounds safe, right? Now imagine that an attacker pads a path with a large number of `'/`'s at the beginning. This could result in future operations being performed on the file `"/"`. If the program appends values in the belief that the result will be safe, the program may be exploitable. Or, the attacker could devise a long filename near the buffer length, so that attempts to append to the filename would silently fail to occur (or only partially occur in ways that may be exploitable).

When using statically-allocated buffers, you really need to consider the length of the source and destination arguments. Sanity checking the input and the resulting intermediate computation might deal with this, too.

Another alternative is to dynamically reallocate all strings instead of using fixed-size buffers. This general approach is recommended by the GNU programming guidelines, since it permits programs to handle arbitrarily-sized inputs (until they run out of memory). Of course, the major problem with dynamically allocated strings is that you may run out of memory. The memory may even be exhausted at some other point in the program than the portion where you're worried about buffer overflows; any memory

allocation can fail. Also, since dynamic reallocation may cause memory to be inefficiently allocated, it is entirely possible to run out of memory even though technically there is enough virtual memory available to the program to continue. In addition, before running out of memory the program will probably use a great deal of virtual memory; this can easily result in “thrashing”, a situation in which the computer spends all its time just shuttling information between the disk and memory (instead of doing useful work). This can have the effect of a denial of service attack. Some rational limits on input size can help here. In general, the program must be designed to fail safely when memory is exhausted if you use dynamically allocated strings.

### 5.2.3. `strncpy` and `strncat`

An alternative, being employed by OpenBSD, is the `strncpy(3)` and `strncat(3)` functions by Miller and de Raadt [Miller 1999]. This is a minimalist, statically-sized buffer approach that provides C string copying and concatenation with a different (and less error-prone) interface. Source and documentation of these functions are available under a newer BSD-style open source license at <ftp://ftp.openbsd.org/pub/OpenBSD/src/lib/libc/string/strncpy.3>.

First, here are their prototypes:

```
size_t strncpy (char *dst, const char *src, size_t size);
size_t strncat (char *dst, const char *src, size_t size);
```

Both `strncpy` and `strncat` take the full size of the destination buffer as a parameter (not the maximum number of characters to be copied) and guarantee to NIL-terminate the result (as long as `size` is larger than 0). Remember that you should include a byte for NIL in the `size`.

The `strncpy` function copies up to `size-1` characters from the NUL-terminated string `src` to `dst`, NIL-terminating the result. The `strncat` function appends the NIL-terminated string `src` to the end of `dst`. It will append at most `size - strlen(dst) - 1` bytes, NIL-terminating the result.



One minor disadvantage of `strncpy(3)` and `strncat(3)` is that they are not, by default, installed in most Unix-like systems. In OpenBSD, they are part of `<string.h>`. This is not that difficult a problem; since they are small functions, you can even include them in your own program's source (at least as an option), and create a small separate package to load them. You can even use `autoconf` to handle this case automatically. If more programs use these functions, it won't be long before these are standard parts of Linux distributions and other Unix-like systems. Also, these functions have been recently added to the "glib" library (I submitted the patch to do this), so using glib will (in the future) make them available. In glib these functions are named `g_strncpy()` and `g_strncat`.

## 5.2.4. libmib

One toolset for C that dynamically reallocates strings automatically is the "libmib allocated string functions" by Forrest J. Cavalier III, available at <http://www.mibsoftware.com/libmib/astring>. There are two variations of libmib; "libmib-open" appears to be clearly open source under its own X11-like license that permits modification and redistribution, but redistributions must choose a different name, however, the developer states that it "may not be fully tested." To continuously get libmib-mature, you must pay for a subscription. The documentation is not open source, but it is freely available.

## 5.2.5. Libsafe

Arash Baratloo, Timothy Tsai, and Navjot Singh (of Lucent Technologies) have developed Libsafe, a wrapper of several library functions known to be vulnerable to stack smashing attacks. This wrapper (which they call a kind of "middleware") is a simple dynamically loaded library that contains modified versions of C library functions such as `strcpy(3)`. These modified versions implement the original functionality, but in a manner that ensures that any buffer overflows are contained within the current stack frame. Their initial performance analysis suggests that this library's overhead is very small. Libsafe papers and source code are available at

<http://www.bell-labs.com/org/11356/libsafe.html>. The Libsafe source code is available under the completely open source LGPL license, and there are indications that many Linux distributors are interested in using it.

Libsafe's approach appears somewhat useful. Libsafe should certainly be considered for inclusion by Linux distributors, and its approach is worth considering by others as well. However, as a software developer, Libsafe is a useful mechanism to support defense-in-depth but it does not really prevent buffer overflows. Here are several reasons why you shouldn't depend just on Libsafe during code development:

- Libsafe only protects a small set of known functions with obvious buffer overflow issues. At the time of this writing, this list is significantly shorter than the list of functions in this paper known to have this problem. It also won't protect against code you write yourself (e.g., in a while loop) that causes buffer overflows.
- Even if libsafe is installed in a distribution, the way it is installed impacts its use. The documentation recommends setting `LD_PRELOAD` to cause libsafe's protections to be enabled, but the problem is that users can unset this environment variable... causing the protection to be disabled for programs they execute!
- Libsafe only protects against buffer overflows of the stack onto the return address; you can still overrun the heap or other variables in that procedure's frame.
- Unless you can be assured that all deployed platforms will use libsafe (or something like it), you'll have to protect your program as though it wasn't there.
- LibSafe seems to assume that saved frame pointers are at the beginning of each stack frame. This isn't always true. Compilers (such as gcc) can optimize away things, and in particular the option `"-fomit-frame-pointer"` removes the information that libsafe seems to need. Thus, libsafe may fail to work for some programs.

The libsafe developers themselves acknowledge that software developers shouldn't just depend on libsafe. In their words:

It is generally accepted that the best solution to buffer overflow attacks is to fix the defective programs. However, fixing defective programs requires knowing that a particular program is defective. The true benefit of using libsafe and other alternative security measures is

protection against future attacks on programs that are not yet known to be vulnerable.

## 5.2.6. Other Libraries

The glib (not glibc) library is a widely-available open source library that provides a number of useful functions for C programmers. GTK+ and GNOME both use glib, for example. As I noted earlier, glib version 2.0 will include `strncpy()` and `strncat()` (through my patch), making it easier to portably use those functions. At this time I do not have an analysis showing definitively that the glib library functions protect against buffer overflow. However, many of the glib functions automatically allocate memory, and those functions automatically *fail with no reasonable way to intercept the failure* (e.g., to try something else instead). As a result, in many cases most glib functions cannot be used in most secure programs. The GNOME guidelines recommend using functions such as `g_strdup_printf()`, which is fine as long as it's okay if your program immediately crashes if an out-of-memory condition occurs. However, if you can't do this, then using such routines isn't appropriate.

## 5.3. Compilation Solutions in C/C++

A completely different approach is to use compilation methods that perform bounds-checking (see [Sitaker 1999] for a list). In my opinion, such tools are very useful in having multiple layers of defense, but it's not wise to use this technique as your sole defense. There are at least two reasons for this. First of all, most such tools only provide partial defense against buffer overflows (and the “complete” defenses are generally 12-30 times slower); C and C++ were simply not designed to protect against buffer overflow. Second of all, for open source programs you cannot be certain what tools will be used to compile the program; using the default “normal” compiler for a given system might suddenly open security flaws.

One of the more useful tools is “StackGuard”, a modification of the standard GNU C compiler `gcc`. StackGuard works by inserting a “guard” value (called a “canary”) in

front of the return address; if a buffer overflow overwrites the return address, the canary's value (hopefully) changes and the system detects this before using it. This is quite valuable, but note that this does not protect against buffer overflows overwriting other values (which they may still be able to use to attack a system). There is work to extend StackGuard to be able to add canaries to other data items, called "PointGuard". PointGuard will automatically protect certain values (e.g., function pointers and longjump buffers). However, protecting other variable types using PointGuard requires specific programmer intervention (the programmer has to identify which data values must be protected with canaries). This can be valuable, but it's easy to accidentally omit protection for a data value you didn't think needed protection - but needs it anyway. More information on StackGuard, PointGuard, and other alternatives is in Cowan [1999].

As a related issue, in Linux you could modify the Linux kernel so that the stack segment is not executable; such a patch to Linux does exist (see Solar Designer's patch, which includes this, at <http://www.openwall.com/linux/> However, as of this writing this is not built into the Linux kernel. Part of the rationale is that this is less protection than it seems; attackers can simply force the system to call other "interesting" locations already in the program (e.g., in its library, the heap, or static data segments). Also, sometimes Linux does require executable code in the stack, e.g., to implement signals and to implement GCC "trampolines". Solar Designer's patch does handle these cases, but this does complicate the patch. Personally, I'd like to see this merged into the main Linux distribution, since it does make attacks somewhat more difficult and it defends against a range of existing attacks. However, I agree with Linus Torvalds and others that this does not add the amount of protection it would appear to and can be circumvented with relative ease. You can read Linus Torvalds' explanation for not including this support at <http://lwn.net/980806/a/linus-noexec.html>.

In short, it's better to work first on developing a correct program that defends itself against buffer overflows. Then, after you've done this, by all means use techniques and tools like StackGuard as an additional safety net. If you've worked hard to eliminate buffer overflows in the code itself, then StackGuard is likely to be more effective because there will be fewer "chinks in the armor" that StackGuard will be called on to protect.

## 5.4. Other Languages

The problem of buffer overflows is an excellent argument for using other programming languages such as Perl, Python, Java, and Ada95. After all, nearly all other programming languages used today (other than assembly language) protect against buffer overflows. Using those other languages does not eliminate all problems, of course; in particular see the discussion in Section 7.1 regarding the NIL character. There is also the problem of ensuring that those other languages' infrastructure (e.g., run-time library) is available and secured. Still, you should certainly consider using other programming languages when developing secure programs to protect against buffer overflows.

# Chapter 6. Structure Program Internals and Approach

*Like a city whose walls are  
broken down is a man who lacks  
self-control.*

*Proverbs 25:28 (NIV)*

## 6.1. Secure the Interface

Interfaces should be minimal (simple as possible), narrow (provide only the functions needed), and non-bypassable. Trust should be minimized. Consider limiting the data that the user can see.

Applications and data viewers may be used to display files developed externally, so in general don't allow them to accept programs (also known as "scripts" or "macros") unless you're willing to do the extensive work necessary to create a secure sandbox. The most dangerous kind is an auto-executing macro that executes when the application is loaded and/or when the data is initially displayed; from a security point-of-view this is a disaster waiting to happen unless you have extremely strong control over what the macro can do (a "sandbox"), and past experience has shown that real sandboxes are hard to implement.

## 6.2. Minimize Privileges

As noted earlier, it is an important general principle that programs have the minimal amount of privileges necessary to do its job (this is termed "least privilege"). That way, if the program is broken, its damage is limited. The most extreme example is to simply not write a secure program at all - if this can be done, it usually should be. For example,

don't make your program `setuid` or `setgid` if you can; just make it an ordinary program, and require the administrator to log in as such before running it.

In Linux and Unix, the primary determiner of a process' privileges is the set of id's associated with it: each process has a real, effective and saved id for both the user and group. Linux also has the filesystem uid and gid. Manipulating these values is critical to keeping privileges minimized, and there are several ways to minimize them (discussed below). You can also use `chroot(2)` to minimize the files visible to a program.

### 6.2.1. Minimize the Privileges Granted

Perhaps the most effective technique is to simply minimize the the highest privilege granted. In particular, avoid granting a program root privilege if possible. Don't make a program `setuid root` if it only needs access to a small set of files; consider creating separate user or group accounts for different function.

A common technique is to create a special group, change a file's group ownership to that group, and then make the program `setgid` to that group. It's better to make a program `setgid` instead of `setuid` where you can, since group membership grants fewer rights (in particular, it does not grant the right to change file permissions).

This is commonly done for game high scores. Games are usually `setgid games`, the score files are owned by the group `games`, and the programs themselves and their configuration files are owned by someone else (say root). Thus, breaking into a game allows the perpetrator to change high scores but doesn't grant the privilege to change the game's executable or configuration file. The latter is important; if an attacker could change a game's executable or its configuration files (which might control what the executable runs), then they might be able to gain control of a user who ran the game.

If creating a new group isn't sufficient, consider creating a new pseudouser (really, a special role) to manage a set of resources. Web servers typically do this; often web servers are set up with a special user ("nobody") so that they can be isolated from other users. Indeed, web servers are instructive here: web servers typically need root privileges to start up (so they can attach to port 80), but once started they usually shed all their privileges and run as the user "nobody". Again, usually the pseudouser doesn't own the primary program it runs, so breaking into the account doesn't allow for

changing the program itself. As a result, breaking into a running web server normally does not automatically break the whole system's security.

If you *must* give a program root privileges, consider using the POSIX capability features available in Linux 2.2 and greater to minimize them immediately on program startup. By calling `cap_set_proc(3)` or the Linux-specific `capsetp(3)` routines immediately after starting, you can permanently reduce the abilities of your program to just those abilities it actually needs. Note that *not* all Unix-like systems implement POSIX capabilities, so this is an approach that can lose portability; however, if you use it merely as an optional safeguard only where it's available, using this approach will not really limit portability. Also, while the Linux kernel version 2.2 and greater includes the low-level calls, the C-level libraries to make their use easy are not installed on some Linux distributions, slightly complicating their use in applications. For more information on Linux's implementation of POSIX capabilities, see <http://linux.kernel.org/pub/linux/libs/security/linux-privs>.

One Linux-unique tool you can use to simplify minimizing granted privileges is the “compartment” tool developed by SuSE. This tool sets the filesystem root, uid, gid, and/or the capability set, then runs the given program. This is particularly handy for running some other program without modifying it. Here's the syntax of version 0.5:

```
Syntax: compartment [options] /full/path/to/program
```

Options:

```
-chroot path    chroot to path
-user user      change uid to this user
-group group    change gid to this group
-init program   execute this program before doing anything
-cap capset     set capset name. You can specify several
-verbose        be verbose
-quiet          do no logging (to syslog)
```

Thus, you could start a more secure anonymous ftp server using:

```
compartment -chroot /home/ftp -cap CAP_NET_BIND_SERVICE anon-ftpd
```



At the time of this writing, the tool is immature and not available on typical Linux distributions, but this may quickly change. You can download the program via <http://www.suse.de/~marc>.

## 6.2.2. Minimize the Time the Privilege Can Be Used

As soon as possible, permanently give up privileges. Some Unix-like systems, including Linux, implement “saved” IDs which store the “previous” value. The simplest approach is to set the other id’s twice to an untrusted id. In `setuid/setgid` programs, you should usually set the effective gid and uid to the real ones, in particular right after a `fork(2)`, unless there’s a good reason not to. Note that you have to change the gid first when dropping from root to another privilege or it won’t work - once you drop root privileges, you won’t be able to change much else.

It’s worth noting that there’s a well-known related bug that uses POSIX capabilities to interfere with this minimization. This bug affects Linux kernel 2.2.0 through 2.2.15, and possibly a number of other Unix-like systems with POSIX capabilities. See Bugtraq id 1322 on <http://www.securityfocus.com> for more information. Here is their summary:

POSIX "Capabilities" have recently been implemented in the Linux kernel. These "Capabilities" are an additional form of privilege control to enable more specific control over what privileged processes can do. Capabilities are implemented as three (fairly large) bitfields, which each bit representing a specific action a privileged process can perform. By setting specific bits, the actions of privileged processes can be controlled – access can be granted for various functions only to the specific parts of a program that require them. It is a security measure. The problem is that capabilities are copied with `fork()` execs, meaning that if capabilities are modified by a parent process, they can be carried over. The way that this can be exploited is by setting all of the capabilities to zero (meaning, all of the bits are off) in each of the three bitfields and then executing a `setuid` program that attempts to drop privileges before executing code that could be dangerous if run as root, such as what `sendmail` does. When `sendmail` attempts to drop privileges using `setuid(getuid())`, it fails not having the capabilities required to do so in its bitfields and with no checks on its return value . It continues executing with superuser privileges, and can run a users `.forward` file as root leading to a complete compromise.

One approach, used by sendmail, is to attempt to do `setuid(0)` after a `setuid(getuid());`; normally this should fail. If it succeeds, the program should stop. For more information, see <http://sendmail.net/?feed=000607linuxbug>. In the short term this might be a good idea in other programs, though clearly the better long-term approach is to upgrade the underlying system.

### 6.2.3. Minimize the Time the Privilege is Active

Use `setuid(2)`, `seteuid(2)`, and related functions to ensure that the program only has these privileges active when necessary. As noted above, you might want ensure that these privileges are disabled while parsing user input, but more generally, only turn on privileges when they're actually needed. Note that some buffer overflow attacks, if successful, can force a program to run arbitrary code, and that code could re-enable privileges that were temporarily dropped. Thus, it's always better to completely drop privileges as soon as possible. Still, temporarily disabling these permissions prevents a whole class of attacks, such as techniques to convince a program to write into a file that perhaps it didn't intent to write into. Since this technique prevents many attacks, it's worth doing if completely dropping the privileges can't be done at that point in the program.

### 6.2.4. Minimize the Modules Granted the Privilege

If only a few modules are granted the privilege, then it's much easier to determine if they're secure. One way to do so is to have a single module use the privilege and then drop it, so that other modules called later cannot misuse the privilege. Another approach is to have separate commands in separate executables; one command might be a complex tool that can do a vast number of tasks for a privileged user (e.g., root), while the other tool is `setuid` but is a small, simple tool that only permits a small command subset. The small, simple tool checks to see if the input meets various criteria for acceptability, and then if it determines the input is acceptable, it passes the input is passed to the tool. This can even be layered several ways, for example, a complex user tool could call a simple `setuid` "wrapping" program (that checks its inputs for secure

values) that then passes on information to another complex trusted tool. This approach is especially helpful for GUI-based systems; have the GUI portion run as a normal user, and then pass security-relevant requests on to another program that has the special privileges for actual execution.

Some operating systems have the concept of multiple layers of trust in a single process, e.g., Multics' rings. Standard Unix and Linux don't have a way of separating multiple levels of trust by function inside a single process like this; a call to the kernel increases privileges, but otherwise a given process has a single level of trust. Linux and other Unix-like systems can sometimes simulate this ability by forking a process into multiple processes, each of which has different privilege. To do this, set up a secure communication channel (usually unnamed pipes or unnamed sockets are used), then fork into different processes and have each process drop as many privileges as possible. Then use a simple protocol to allow the less trusted processes to request actions from the more trusted process(es), and ensure that the more trusted processes only support a limited set of requests.

This is one area where technologies like Java 2 and Fluke have an advantage. For example, Java 2 can specify fine-grained permissions such as the permission to only open a specific file. However, general-purpose operating systems do not typically have such abilities at this time; this may change in the near future.

### **6.2.5. Consider Using FSUID To Limit Privileges**

Each Linux process has two Linux-unique state values called filesystem user id (fsuid) and filesystem group id (fsgid). These values are used when checking against the filesystem permissions. If you're building a program that operates as a file server for arbitrary users (like an NFS server), you might consider using these Linux extensions. To use them, while holding root privileges change just fsuid and fsgid before accessing files on behalf of a normal user. This extension is fairly useful, and provides a mechanism for limiting filesystem access rights without removing other (possibly necessary) rights. By only setting the fsuid (and not the euid), a local user cannot send a signal to the process. Also, avoiding race conditions is much easier in this situation. However, a disadvantage of this approach is that these calls are not portable to other

Unix-like systems.

## 6.2.6. Consider Using Chroot to Minimize Available Files

You can use `chroot(2)` to limit the files visible to your program. This requires carefully setting up a directory (called the “chroot jail”) and correctly entering it. This can be a fairly effective technique for improving a program’s security - it’s hard to interfere with files you can’t see. However, it depends on a whole bunch of assumptions, in particular, the program must lack root privileges, it must not have any way to get root privileges, and the chroot jail must be properly set up. I recommend using `chroot(2)` where it makes sense to do so, but don’t depend on it alone; instead, make it part of a layered set of defenses. Here are a few notes about the use of `chroot(2)`:

- The program can still use non-filesystem objects that are shared across the entire machine (such as System V IPC objects and network sockets). It’s best to also use separate pseudousers and/or groups, because all Unix-like systems include the ability to isolate users; this will at least limit the damage a subverted program can do to other programs. Note that current most Unix-like systems (including Linux) won’t isolate intentionally cooperating programs; if you’re worried about malicious programs cooperating, you need to get a system that implements some sort of mandatory access control and/or limits covert channels.
- Be sure to close any filesystem descriptors to outside files if you don’t want them used later. In particular, don’t have any descriptors open to directories outside the chroot jail, or set up a situation where such a descriptor could be given to it (e.g., via Unix sockets or an old implementation of `/proc`). If the program is given a descriptor to a directory outside the chroot jail, it could be used to escape out of the chroot jail.
- The chroot jail has to be set up to be secure. Don’t use a normal user’s home directory (or subdirectory) as a chroot jail; use a separate location or “home” directory specially set aside for the purpose. Place the absolute minimum number of files there. Typically you’ll have a `/bin`, `/etc/`, `/lib`, and maybe one or two others (e.g., `/pub` if it’s an ftp server). Place in `/bin` only what you need to run after doing the

chroot()); sometimes you need nothing at all (try to avoid placing a shell there, though sometimes that can't be helped). You may need a `/etc/passwd` and `/etc/group` so file listings can show some correct names, but if so, try not to include the real system's values, and certainly replace all passwords with `"*"`. In `/lib`, place only what you need; use `ldd(1)` to query each program in `/bin` to find out what it needs, and only include them. On Linux, you'll probably need a few basic libraries like `ld-linux.so.2`, and not much else. It's usually wiser to completely copy in all files, instead of making hard links; while this wastes some time and disk space, it makes it so that attacks on the chroot jail files do not automatically propagate into the regular system's files. Mounting a `/proc` filesystem, on systems where this is supported, is generally unwise. In fact, in 2.0.x versions of Linux it's a known security flaw, since there are pseudodirectories in `/proc` that would permit a chroot'ed program to escape. Linux kernel 2.2 fixed this known problem, but there may be others; if possible, don't do it.

- Chroot really isn't effective if the program can acquire root privilege. For example, the program could use calls like `mknod(2)` to create a device file that can view physical memory, and then use the resulting device file to modify kernel memory to give itself whatever privileges it desired. Another example of how a root program can break out of chroot is demonstrated at <http://www.suid.edu/source/breakchroot.c>. In this example, the program opens a file descriptor for the current directory, creates and chroots into a subdirectory, sets the current directory to the previously-opened current directory, repeatedly `cd`'s up from the current directory (which since it is outside the current chroot succeeds in moving up to the real filesystem root), and then calls `chroot` on the result. By the time you read this, these weaknesses may have been plugged, but the reality is that root privilege has traditionally meant "all privileges" and it's hard to strip them away. It's better to assume that a program requiring continuous root privileges will only be mildly helped using `chroot()`. Of course, you may be able to break your program into parts, so that at least part of it can be in a chroot jail.

### 6.2.7. Consider Minimizing the Accessible Data

Consider minimizing the amount of data that can be accessed by the user. For example, in CGI scripts, place all data used by the CGI script outside of the document tree unless there is a reason the user needs to see the data directly. Some people have the false notion that, by not publically providing a link, no one can access the data, but this is simply not true.

### **6.3. Avoid Creating Setuid/Setgid Scripts**

Many Unix-like systems, in particular Linux, simply ignore the setuid and setgid bits on scripts to avoid the race condition described earlier. Since support for setuid scripts varies on Unix-like systems, they're best avoided in new applications where possible. As a special case, Perl includes a special setup to support setuid Perl scripts, so using setuid and setgid is acceptable in Perl if you truly need this kind of functionality. If you need to support this kind of functionality in your own interpreter, examine how Perl does this. Otherwise, a simple approach is to “wrap” the script with a small setuid/setgid executable that creates a safe environment (e.g., clears and sets environment variables) and then calls the script (using the script's full path). Make sure that the script cannot be changed by an attacker! Shell scripting languages have additional problems, and really should not be setuid/setgid; see Section 9.4 for more information about this.

### **6.4. Configure Safely and Use Safe Defaults**

Configuration is considered to currently be the number one security problem. Therefore, you should spend some effort to (1) make the initial installation secure, and (2) make it easy to reconfigure the system while keeping it secure.

A program should have the most restrictive access policy until the administrator has a chance to configure it. Please don't create “sample” working users or “allow access to all” configurations as the starting configuration; many users just “install everything”

(installing all available services) and never get around to configuring many services. In some cases the program may be able to determine that a more generous policy is reasonable by depending on the existing authentication system, for example, an ftp server could legitimately determine that a user who can log into a user's directory should be allowed to access that user's files. Be careful with such assumptions, however.

Have installation scripts install a program as safely as possible. By default, install all files as owned by root or some other system user and make them unwriteable by others; this prevents non-root users from installing viruses. Indeed, it's best to make them unreadable by all but the trusted user. Allow non-root installation where possible as well, so that users without root privileges and administrators who do not fully trust the installer can still use the program.

Try to make configuration as easy and clear as possible, including post-installation configuration. Make using the "secure" approach as easy as possible, or many users will use an insecure approach without understanding the risks. On Linux, take advantage of tools like `linuxconf`, so that users can easily configure their system using an existing infrastructure.

If there's a configuration language, the default should be to deny access until the user specifically grants it. Include many clear comments in the sample configuration file, if there is one, so the administrator understands what the configuration does.

## 6.5. Fail Safe

A secure program should always "fail safe", that is, it should be designed so that if the program does fail, the safest result should occur. For security-critical programs, that usually means that if some sort of misbehavior is detected (malformed input, reaching a "can't get here" state, and so on), then the program should immediately deny service and stop processing that request. Don't try to "figure out what the user wanted": just deny the service. Sometimes this can decrease reliability or usability (from a user's perspective), but it increases security. There are a few cases where this might not be desired (e.g., where denial of service is much worse than loss of confidentiality or

integrity), but such cases are quite rare.

Note that I recommend “stop processing the request”, not “fail altogether”. In particular, most servers should not completely halt when given malformed input, because that creates a trivial opportunity for a denial of service attack (the attacker just sends garbage bits to prevent you from using the service). Sometimes taking the whole server down is necessary, in particular, reaching some “can’t get here” states may signal a problem so drastic that continuing is unwise.

Consider carefully what error message you send back when a failure is detected. If you send nothing back, it may be hard to diagnose problems, but sending back too much information may unintentionally aid an attacker. Usually the best approach is to reply with “access denied” or “miscellaneous error encountered” and then write more detailed information to an audit log (where you can have more control over who sees the information).

## 6.6. Avoid Race Conditions

A “race condition” can be defined as “Anomalous behavior due to unexpected critical dependence on the relative timing of events” [FOLDOC]. Race conditions generally involve one or more processes accessing a shared resource (such a file or variable), where this multiple access has not been properly controlled.

In general, processes do not execute atomically; another process may interrupt it between essentially any two instructions. If a secure program’s process is not prepared for these interruptions, another process may be able to interfere with the secure program’s process. Any pair of operations must not fail if another process’s code arbitrary code is executed between them.

Race condition problems can be notionally divided into two categories:

- Interference caused by untrusted processes. Some security taxonomies call this problem a “sequence” or “non-atomic” condition. These are conditions caused by processes running other, different programs, which “slip in” other actions between



steps of the secure program. These other programs might be invoked by an attacker specifically to cause the problem. This paper will call these sequencing problems.

- Interference caused by trusted processes (from the secure program's point of view). Some taxonomies call these deadlock, livelock, or locking failure conditions. These are conditions caused by processes running the "same" program. Since these different processes may have the "same" privileges, if not properly controlled they may be able to interfere with each other in a way other programs can't. Sometimes this kind of interference can be exploited. This paper will call these locking problems.

### 6.6.1. Sequencing (Non-Atomic) Problems

In general, you must check your code for any pair of operations that might fail if arbitrary code is executed between them.

Note that loading and saving a shared variable are usually implemented as separate operations and are not atomic. This means that an "increment variable" operation is usually converted into loading, incrementing, and saving operation, so if the variable memory is shared the other process may interfere with the incrementing.

Secure programs must determine if a request should be granted, and if so, act on that request. There must be no way for an untrusted user to change anything used in this determination before the program acts on it. This kind of race condition is sometimes termed a "time of check - time of use" (TOCTOU) race condition.

This issue repeatedly comes up in the filesystem. Programs should generally avoid using `access(2)` to determine if a request should be granted, followed later by `open(2)`, because users may be able to move files around between these calls, possibly creating symbolic links or files of their own choosing instead. A secure program should instead set its effective id or filesystem id, then make the open call directly. It's possible to use `access(2)` securely, but only when a user cannot affect the file or any directory along its path from the filesystem root.

For example, when performing a series of operations on a file's metainformation (such as changing its owner, stat-ing the file, or changing its permission bits), first open the

file and then use the operations on open files. This means use the `fchown()`, `fstat()`, or `fchmod()` system calls, instead of the functions taking filenames such as `chown()`, `chgrp()`, and `chmod()`. Doing so will prevent the file from being replaced while your program is running (a possible race condition). For example, if you close a file and then use `chmod()` to change its permissions, an attacker may be able to remove the file between those two steps and create a symbolic link to another file (say `/etc/passwd`). Other interesting files include `/dev/zero`, which can provide an infinitely-long data stream of input to a program. Also, avoid the use of the `access()` function to determine your ability to access a file: using the `access()` function followed by an `open()` is a race condition, and almost always a bug. This is only necessary if it's possible for an untrusted process to modify the relevant directory its ancestors.

This issue particularly comes up in the `/tmp` and `/var/tmp` directories, which are shared by all users. Avoid using these directories and their subdirectories if possible. In particular, imagine what would happen if users created files (including symbolic links) at arbitrary times in directories you intend to use (for example, between the time you compute a filename and the time you try to open it). A common trick is for attackers to create symbolic links in the temporary directory to some other file (e.g., `/etc/passwd`) while your secure program is running. The goal is to create a situation where the secure program determines that a given filename doesn't exist, the attacker then creates the symbolic link to another file, and then the secure program opens the file (but now it actually opened an unintended file). Often important files can be clobbered or modified this way.

The general problem when creating files in these shared directories is that you must guarantee that the filename you plan to use doesn't already exist at time of creation (checking "before" you create it doesn't work). Using an "unpredictable" or "unique" filename doesn't work, because another process can often repeatedly guess what that value will be. The GNOME programming guidelines recommend the following C code when creating filesystem objects in shared (temporary) directories to counteract this problem [Quintero 2000]:

```
char *filename;  
int fd;  
  
do {
```

```
filename = tempnam (NULL, "foo");
fd = open (filename, O_CREAT | O_EXCL | O_TRUNC | O_RDWR, 0600);
free (filename);
} while (fd == -1);
```

Note that you need to free() the filename. You should close() and unlink() the file after you are done. If you want to use the Standard C I/O library, you can use fdopen() to transform the file descriptor into a FILE \*. Note that this won't work over NFS version 2 (v2) systems, because older NFS doesn't correctly support O\_EXCL.

## 6.6.2. Locking

There are often situations in which a program must ensure that it has exclusive rights to something (e.g., a file, a device, and/or existence of a particular server process). Any system which locks resources must deal with the standard problems of locks, namely, deadlocks (“deadly embraces”), livelocks, and releasing “stuck” locks if a program doesn't clean up its locks. A deadlock can occur if programs are stuck waiting for each other to release resources. For example, a deadlock would occur if process 1 locks resources A and waits for resource B, while process 2 locks resource B and waits for resource A. Many deadlocks can be prevented by simply requiring all processes that lock multiple resources to lock them in the same order (e.g., alphabetically by lock name).

### 6.6.2.1. Using Files as Locks

On Unix-like systems resource locking has traditionally been done by creating a file to indicate a lock, because this is very portable. It also makes it easy to “fix” stuck locks, because an administrator can just look at the filesystem to see what locks have been set. Stuck locks can occur because the program failed to clean up after itself (e.g., it crashed or malfunctioned) or because the whole system crashed. Note that these are “advisory” (not “mandatory”) locks - all processes needed the resource must cooperate to use these locks.

However, there are several traps to avoid. First, don't use the technique used by very old Unix C programs, which is calling `creat()` or its `open()` equivalent, the `open()` mode `O_WRONLY | O_CREAT | O_TRUNC`, with the file mode set to 0 (no permissions). For normal users on normal file systems, this works, but this approach fails to lock the file when the user has root privileges. Root can always perform this operation, even when the file already exists. In fact, old versions of Unix had this particular problem in the old editor "ed" – the symptom was that occasionally portions of the password file would be placed in user's files! [Rochkind 1985, 22]. Instead, if you're creating a lock for processes that are on the local filesystem, you should use `open()` with the flags `O_WRONLY | O_CREAT | O_EXCL` (and again, no permissions, so that other processes with the same owner won't get the lock). Note the use of `O_EXCL`, which is the official way to create "exclusive" files; this even works for root on a local filesystem. [Rochkind 1985, 27].

Second, if the lock file may be on an NFS-mounted filesystem, then you have the problem that NFS version 2 doesn't completely support normal file semantics. This can even be a problem for work that's supposed to be "local" to a client, since some clients don't have local disks and may have *all* files remotely mounted via NFS. The manual for `open(2)` explains how to handle things in this case (which also handles the case of root programs):

"... programs which rely on [the `O_CREAT` and `O_EXCL` flags of `open(2)`] for performing locking tasks will contain a race condition. The solution for performing atomic file locking using a lockfile is to create a unique file on the same filesystem (e.g., incorporating hostname and pid), use `link(2)` to make a link to the lockfile and use `stat(2)` on the unique file to check if its link count has increased to 2. Do not use the return value of the `link(2)` call."

Obviously, this solution only works if all programs doing the locking are cooperating, and if all non-cooperating programs aren't allowed to interfere. In particular, the directories you're using for file locking must not have permissive file permissions for creating and removing files.

NFS version 3 added support for `O_EXCL` mode in `open(2)`; see IETF RFC 1813, in particular the "EXCLUSIVE" value to the "mode" argument of "CREATE". Sadly, not everyone has switched to NFS version 3 or higher at the time of this writing, so you

you can't depend on this yet in portable programs. Still, in the long run there's hope that this issue will go away.

If you're locking a device or the existence of a process on a local machine, try to use standard conventions. I recommend using the Filesystem Hierarchy Standard (FHS); it is widely referenced by Linux systems, but it also tries to incorporate the ideas of other Unix-like systems. The FHS describes standard conventions for such locking files, including naming, placement, and standard contents of these files [FHS 1997]. If you just want to be sure that your server doesn't execute more than once on a given machine, you should usually create a process identifier as `/var/run/NAME.pid` with the pid as its contents. In a similar vein, you should place lock files for things like device lock files in `/var/lock`. This approach has the minor disadvantage of leaving files hanging around if the program suddenly halts, but it's standard practice and that problem is easily handled by other system tools.

It's important that the programs which are cooperating using files to represent the locks use the same directory, not just the same directory name. This is an issue with networked systems: the FHS explicitly notes that `/var/run` and `/var/lock` are unshareable, while `/var/mail` is shareable. Thus, if you want the lock to work on a single machine, but not interfere with other machines, use unshareable directories like `/var/run` (e.g., you want to permit each machine to run its own server). However, if you want all machines sharing files in a network to obey the lock, you need to use a directory that they're sharing; `/var/mail` is one such location. See FHS section 2 for more information on this subject.

### 6.6.2.2. Other Approaches to Locking

Of course, you need not use files to represent locks. Network servers often need not bother; the mere act of binding to a port acts as a kind of lock, since if there's an existing server bound to a given port, no other server will be able to bind to that port.

Another approach to locking is to use POSIX record locks, implemented through `fcntl(2)` as a "discretionary lock". These are discretionary, that is, using them requires the cooperation of the programs needing the locks (just as the approach to using files to represent locks does). There's a lot to recommend POSIX record locks: POSIX record

locking is supported on nearly all Unix-like platforms (it's mandated by POSIX.1), it can lock portions of a file (not just a whole file), and it can handle the difference between read locks and write locks. Even more usefully, if a process dies, its locks are automatically removed, which is usually what is desired.

You can also use mandatory locks, which are based on System V's mandatory locking scheme. These only apply to files where the locked file's setgid bit is set, but the group execute bit is not set. Also, you must mount the filesystem to permit mandatory file locks. In this case, every `read(2)` and `write(2)` is checked for locking; while this is more thorough than advisory locks, it's also slower. Also, mandatory locks don't port as widely to other Unix-like systems (they're available on Linux and System V-based systems, but not necessarily on others). Note that processes with root privileges can be held up by a mandatory lock, too, making it possible that this could be the basis of a denial-of-service attack.

## 6.7. Trust Only Trustworthy Channels

In general, do not trust results from untrustworthy channels.

In most computer networks (and certainly for the Internet at large), no unauthenticated transmission is trustworthy. For example, on the Internet arbitrary packets can be forged, including header values, so don't use their values as your primary criteria for security decisions unless you can authenticate them. In some cases you can assert that a packet claiming to come from the "inside" actually does, since the local firewall would prevent such spoofs from outside, but broken firewalls, alternative paths, and mobile code make even this assumption suspect. In a similar vein, do not assume that low port numbers (less than 1024) are trustworthy; in most networks such requests can be forged or the platform can be made to permit use of low-numbered ports.

If you're implementing a standard and inherently insecure protocol (e.g., ftp and rlogin), provide safe defaults and document clearly the assumptions.

The Domain Name Server (DNS) is widely used on the Internet to maintain mappings

between the names of computers and their IP (numeric) addresses. The technique called “reverse DNS” eliminates some simple spoofing attacks, and is useful for determining a host’s name. However, this technique is not trustworthy for authentication decisions. The problem is that, in the end, a DNS request will be sent eventually to some remote system that may be controlled by an attacker. Therefore, treat DNS results as an input that needs validation and don’t trust it for serious access control.

If asking for a password, try to set up trusted path (e.g., require pressing an unforgeable key before login, or display unforgeable pattern such as flashing LEDs). Otherwise, an “evil” program could create a display that “looks like” the expected display for a password (e.g., a log-in) and intercept that password. Unfortunately, stock Linux and most other Unixes don’t have a trusted path even for its normal login sequence, and since currently normal users can change the LEDs, the LEDs can’t currently be used to confirm a trusted path. When handling a password over a network, encrypt it between trusted endpoints.

Arbitrary email (including the “from” value of addresses) can be forged as well. Using digital signatures is a method to thwart many such attacks. A more easily thwarted approach is to require emailing back and forth with special randomly-created values, but for low-value transactions such as signing onto a public mailing list this is usually acceptable.

If you need a trustworthy channel over an untrusted network, you need some sort of cryptologic service (at the very least, a cryptologically safe hash); see Section 10.4 for more information on cryptographic algorithms and protocols.

Note that in any client/server model, including CGI, that the server must assume that the client can modify any value. For example, so-called “hidden fields” and cookie values can be changed by the client before being received by CGI programs. These cannot be trusted unless special precautions are taken. For example, the hidden fields could be signed in a way the client cannot forge as long as the server checks the signature. The hidden fields could also be encrypted using a key only the trusted server could decrypt (this latter approach is the basic idea behind the Kerberos authentication system). InfoSec labs has further discussion about hidden fields and applying encryption at <http://www.infoseclabs.com/mschff/mschff.htm>. In general, you’re better off keeping data you care about at the server end in a client/server model. In the same

vein, don't depend on `HTTP_REFERER` for authentication in a CGI program, because this is sent by the user's browser (not the web server).

The routines `getlogin(3)` and `ttyname(3)` return information that can be controlled by a local user, so don't trust them for security purposes.

This issue applies to data referencing other data, too. For example, HTML or XML allow you to include by reference other files (e.g., DTDs and style sheets) that may be stored remotely. However, those external references could be modified so that users see a very different document than intended; a style sheet could be modified to "white out" words at critical locations, deface its appearance, or insert new text. External DTDs could be modified to prevent use of the document (by adding declarations that break validation) or insert different text into documents [St. Laurent 2000].

## 6.8. Use Internal Consistency-Checking Code

The program should check to ensure that its call arguments and basic state assumptions are valid. In C, macros such as `assert(3)` may be helpful in doing so.

## 6.9. Self-limit Resources

In network daemons, shed or limit excessive loads. Set limit values (using `setrlimit(2)`) to limit the resources that will be used. At the least, use `setrlimit(2)` to disable creation of "core" files. For example, by default Linux will create a core file that saves all program memory if the program fails abnormally, but such a file might include passwords or other sensitive data.



# Chapter 7. Carefully Call Out to Other Resources

*Do not put your trust in princes,  
in mortal men, who cannot save.  
Psalms 146:3 (NIV)*

## 7.1. Limit Call-outs to Valid Values

Ensure that any call out to another program only permits valid and expected values for every parameter. This is more difficult than it sounds, because there are many library calls or commands call lower-level routines in potentially surprising ways. For example, several system calls, such as `popen(3)` and `system(3)`, are implemented by calling the command shell, meaning that they will be affected by shell metacharacters. Similarly, `execlp(3)` and `execvp(3)` may cause the shell to be called. Many guidelines suggest avoiding `popen(3)`, `system(3)`, `execlp(3)`, and `execvp(3)` entirely and use `execve(3)` directly in C when trying to spawn a process [Galvin 1998b]. At the least, avoid using `system(3)` when you can use the `execve(3)`; since `system(3)` uses the shell to expand characters, there is more opportunity for mischief in `system(3)`. In a similar manner the Perl and shell backtick (```) also call a command shell; for more information on Perl see Section 9.2.

One of the nastiest examples of this problem are shell metacharacters. The standard Unix-like command shell (stored in `/bin/sh`) interprets a number of characters specially. If these characters are sent to the shell, then their special interpretation will be used unless escaped; this fact can be used to break programs. According to the WWW Security FAQ [Stein 1999, Q37], these metacharacters are:

```
& ; ` ' \ " | * ? ~ < > ^ ( ) [ ] { } $ \n \r
```

Unfortunately, in real life this isn't a complete list. Here are some other characters that can be problematic:

- `'!` means “not” in an expression (as it does in C); if the return value of a program is tested, prepending `!` could fool a script into thinking something had failed when it succeeded or vice versa. In some shells, the `!"` also accesses the command history, which can cause real problems. In bash, this only occurs for interactive mode, but `tcsh` (a `csh` clone found in some Linux distributions) uses `!"` even in scripts. In `csh`, bash, and some other shells, if you can fool them i Also new bash seems to use `'!` for accessing command history - but this probably only in interactive mode.
- `'#'` is the comment character; all further text is ignored.
- `'-'` can be misinterpreted as leading an option (or, as `-`, disabling all further options). Even if it's in the “middle” of a filename, if it's preceded by what the shell considers as whitespace you may have a problem.
- `' '` (space) and other whitespace characters may turn a “single” filename into multiple arguments.
- Other control characters (in particular, `NIL`) may cause problems for some shell implementations.
- Depending on your usage, it's even conceivable that `“.”` (the “run in current shell”) and `“=”` (for setting variables) might be worrisome characters. However, any example I've found so far where these are issues have other (much worse) security problems.

Forgetting one of these characters can be disastrous, for example, many programs omit backslash as a metacharacter [rfp 1999]. As discussed in the Chapter 4, a recommended approach by some is to immediately escape at least all of these characters when they are input. But again, by far and away the best approach is to identify which characters you wish to permit, and use a filter to only permit those characters.

A number of programs, especially those designed for human interaction, have “escape” codes that perform “extra” activities. One of the more common (and dangerous) escape codes is one that brings up a command line. Make sure that these “escape” commands can't be included (unless you're sure that the specific command is safe). For example,

many line-oriented mail programs (such as mail or mailx) use tilde (~) as an escape character, which can then be used to send a number of commands. As a result, apparently-innocent commands such as “mail admin < file-from-user” can be used to execute arbitrary programs. Interactive programs such as vi, emacs, and ed have “escape” mechanisms that allow users to run arbitrary shell commands from their session. Always examine the documentation of programs you call to search for escape mechanisms.

The issue of avoiding escape codes even goes down to low-level hardware components and emulators of them. Most modems implement the so-called “Hayes” command set, in which the sequence “+++”, a delay, and then “+++” again forces the modem to switch modes (and interpret following text as commands to it). This can be used to implement denial-of-service attacks or even forcing a user to connect to someone else.

Many “terminal” interfaces implement the escape codes of ancient, long-gone physical terminals like the VT100. These codes can be useful, for example, for bolding characters, changing font color, or moving to a particular location in a terminal interface. However, do not allow arbitrary untrusted data to be sent directly to a terminal screen, because some of those codes can cause serious problems. On some systems you can remap keys (e.g., so when a user presses “Enter” or a function key it sends the command you want them to run). On some you can even send codes to clear the screen, display a set of commands you’d like the victim to run, and then send that set “back”, forcing the victim to run the commands of the attacker’s choosing without even waiting for a keystroke. This is typically implemented using “page-mode buffering”. This security problem is why emulated tty’s (represented as device files, usually in /dev/) should only be writeable by their owners and never anyone else - they should never have “other write” permission set, and unless only the user is a member of the group (i.e., the “user-private group” scheme), the “group write” permission should not be set either for the terminal [Filipski 1986]. If you’re displaying data to the user at a (simulated) terminal, you probably need to filter out all control characters (characters with values less than 32) from data sent back to the user unless they’re identified by you as safe. Worse comes to worse, you can identify tab and newline (and maybe carriage return) as safe, removing all the rest. Characters with their high bits set (i.e., values greater than 127) are in some ways trickier to handle; some old systems implement them as if they weren’t set, but simply filtering them inhibits much

international use. In this case, you need to look at the specifics of your situation.

A related problem is that the NIL character (character 0) can have surprising effects. Most C and C++ functions assume that this character marks the end of a string, but string-handling routines in other languages (such as Perl and Ada95) can handle strings containing NIL. Since many libraries and kernel calls use the C convention, the result is that what is checked is not what is actually used [rfp 1999].

When calling another program or referring to a file always specify its full path (e.g., `/usr/bin/sort`). For program calls, this will eliminate possible errors in calling the “wrong” command, even if the PATH value is incorrectly set. For other file referents, this reduces problems from “bad” starting directories.

## 7.2. Call Only Interfaces Intended for Programmers

Call only application programming interfaces (APIs) that are intended for use by programs. Usually a program can invoke any other program, including those that are really designed for human interaction. However, it’s usually unwise to invoke a program intended for human interaction in the same way a human would. The problem is that programs’s human interfaces are intentionally rich in functionality and are often difficult to completely control. Usually the same program can be invoked As discussed in Section 7.1, interactive programs often have “escape” codes; programs shouldn’t usually just call `mail`, `mailx`, `ed`, `vi`, and `emacs`, at least not without checking their input. Usually there are parameters to give you safer access to the program’s functionality, or a different API or application that’s intended for use by programs. Use those instead. For example, instead of an editor (like `ed`, `vi`, or `emacs`), use `sed` where you can.

## 7.3. Check All System Call Returns

Every system call that can return an error condition must have that error condition

checked. One reason is that nearly all system calls require limited system resources, and users can often affect resources in a variety of ways. Setuid/setgid programs can have limits set on them through calls such as `setrlimit(3)` and `nice(2)`. External users of server programs and CGI scripts may be able to cause resource exhaustion simply by making a large number of simultaneous requests. If the error cannot be handled gracefully, then fail open as discussed earlier.

## 7.4. Avoid Using `vfork(2)`

The portable way to create new processes in Unix-like systems is to use the `fork(2)` call. BSD introduced a variant called `vfork(2)` as an optimization technique. In `vfork(2)`, unlike `fork(2)`, the child borrows the parent's memory and thread of control until a call to `execve(2V)` or an exit occurs; the parent process is suspended while the child is using its resources. The rationale is that in old BSD systems, `fork(2)` would actually cause memory to be copied while `vfork(2)` would not. Linux never had this problem; because Linux used copy-on-write semantics internally, Linux only copies pages when they changed (actually, there are still some tables that have to be copied; in most circumstances their overhead is not significant). Nevertheless, since some programs depend on `vfork(2)`, recently Linux implemented the BSD `vfork(2)` semantics (previously `vfork(2)` had been an alias for `fork(2)`).

There are a number of problems with `vfork(2)`. From a portability point-of-view, the problem with `vfork(2)` is that it's actually fairly tricky for a process to not interfere with its parent, especially in high-level languages. The "not interfering" requirement applies to the actual machine code generated, and many compilers generate hidden temporaries and other code structures that cause unintended interference. The result: programs using `vfork(2)` can easily fail when the code changes or even when compiler versions change.

For secure programs it gets worse on Linux systems, because Linux (at least 2.2 versions through 2.2.17) is vulnerable to a race condition in `vfork()`'s implementation. If a privileged process uses a `vfork(2)/execve(2)` pair in Linux to execute user commands, there's a race condition while the child process is already running as the target user's UID, but hasn't entered `execve(2)` yet. The user may be able to send

signals, including SIGSTOP, to this process. Due to the semantics of `vfork(2)`, the privileged parent process would then be blocked as well. As a result, an unprivileged process could cause the privileged process to halt, resulting in a denial-of-service of the privileged process' service. FreeBSD and OpenBSD, at least, have code to specifically deal with this case, so to my knowledge they are not vulnerable to this problem. My thanks to Solar Designer, who noted and documented this problem in Linux on the "security-audit" mailing list on October 7, 2000.

The bottom line with `vfork(2)` is simple: *don't* use `vfork(2)` in your programs. This shouldn't be difficult; the primary use of `vfork(2)` is to support old programs that needed `vfork`'s semantics.

## 7.5. Counter Web Bugs When Retrieving Embedded Content

Some data formats can embed references to content that is automatically retrieved when the data is viewed (not waiting for a user to select it). If it's possible to cause this data to be retrieved through the Internet (e.g., through the World Wide Wide), then there is a potential to use this capability to obtain information about readers without the readers' knowledge, and in some cases to force the reader to perform activities without the reader's consent. This privacy concern is sometimes called a "web bug."

In a web bug, a reference is intentionally inserted into a document and used by the content author to track where (and how often) a document is being read. The author can also watch how a "bugged" document is passed from one person to another or from one organization to another.

The HTML format has had this issue for some time. According to the Privacy Foundation (<http://www.privacyfoundation.org>):

Web bugs are used extensively today by Internet advertising companies on Web pages and in HTML-based email messages for tracking. They are typically 1-by-1 pixel in size to make them invisible on the screen to disguise the fact that they are used for tracking.

What is more concerning is that other document formats seem to have such a capability, too. When viewing HTML from a web site with a web browser, there are other ways of getting information on who is browsing the data, but when viewing a document in another format from an email few users expect that the mere act of reading the document can be monitored. However, for many formats, reading a document can be monitored. For example, it has been recently determined that Microsoft Word can support web bugs; see the Privacy Foundation advisory for more information (<http://www.privacyfoundation.org/advisories/advWordBugs.html>). As noted in their advisory, recent versions of Microsoft Excel and Microsoft Power Point can also be bugged. In some cases, cookies can be used to obtain even more information.

Web bugs are primarily an issue with the design of the file format. If your users value their privacy, you probably will want to limit the automatic downloading of included files. One exception might be when the file itself is being downloaded (say, via a web browser); downloading other files from the same location at the same time is much less likely to be a concern from users.

# Chapter 8. Send Information Back Judiciously

*Do not answer a fool according to his folly, or you will be like him yourself.*

*Proverbs 26:4 (NIV)*

## 8.1. Minimize Feedback

Avoid giving much information to untrusted users; simply succeed or fail, and if it fails just say it failed and minimize information on why it failed. Save the detailed information for audit trail logs. For example:

- If your program requires some sort of user authentication (e.g., you're writing a network service or login program), give the user as little information as possible before they authenticate. In particular, avoid giving away the version number of your program before authentication. Otherwise, if a particular version of your program is found to have a vulnerability, then users who don't upgrade from that version advertise to attackers that they are vulnerable.
- If your program accepts a password, don't echo it back; this creates another way passwords can be seen.

## 8.2. Handle Full/Unresponsive Output

It may be possible for a user to clog or make unresponsive a secure program's output channel back to that user. For example, a web browser could be intentionally halted or



have its TCP/IP channel response slowed. The secure program should handle such cases, in particular it should release locks quickly (preferably before replying) so that this will not create an opportunity for a Denial-of-Service attack. Always place timeouts on outgoing network-oriented write requests.

## 8.3. Control Data Formatting

A number of output routines in computer languages have a parameter that controls the generated format. In C, the most obvious example is the `printf()` family of routines (including `printf()`, `sprintf()`, `snprintf()`, `fprintf()`, and so on). Other examples in C include `syslog()` (which writes system log information) and `setproctitle()` (which sets the string used to display process identifier information). Many functions with names beginning with “err” or “warn”, containing “log”, or ending in “printf” are worth considering. Python includes the “%” operation, which on strings controls formatting in a similar manner. Many programs and libraries define formatting functions, often by calling built-in routines and doing additional processing (e.g., `glib`’s `g_snprintf()` routine).

Surprisingly, many people seem to forget the power of these formatting capabilities and use data from untrusted users as the formatting parameter. Never use unfiltered data from an untrusted user as the format parameter. Perhaps this is best shown by example:

```
/* Wrong ways: */
printf(string_from_untrusted_user);
/* Right ways: */
printf("%s %d", string_from_untrusted_user); /* or just */
fputs(string_from_untrusted_user);
```

Otherwise, an attacker can cause all sorts of mischief by carefully selecting the formatting string. The case of C’s `printf()` is a good example - there are lots of ways to possibly exploit user-controlled format strings in `printf()`. These include buffer overruns by creating a long formatting string (this can result in the attacker having complete control over the program), conversion specifications that use unpassed parameters (causing unexpected data to be inserted), and creating formats which produce totally

unanticipated result values (say by prepending or appending awkward data, causing problems in later use). A particularly nasty case is `printf`'s `%n` conversion specification, which writes the number of characters written so far into the pointer argument; using this, an attacker can overwrite a value that was intended for printing! An attacker can even overwrite almost arbitrary locations, since the attacker can specify a "parameter" that wasn't actually passed. Since in many cases the results are sent back to the user, this attack can also be used to expose internal information about the stack. This information can then be used to circumvent stack protection systems such as StackGuard; StackGuard uses constant "canary" values to detect attacks, but if the stack's contents can be displayed, the current value of the canary will be exposed and made vulnerable.

A formatting string should almost always be a constant string, possibly involving a function call to implement a lookup for internationalization (e.g., via `gettext`'s `_()`). Note that this lookup must be limited to values that the program controls, i.e., the user must be allowed to only select from the message files controlled by the program. It's possible to filter user data before using it (e.g., by designing a filter listing legal characters for the format string such as `[A-Za-z0-9]`), but it's usually better to simply prevent the problem by using a constant format string or `fputs()` instead. Note that although I've listed this as an "output" problem, this can cause problems internally to a program before output (since the output routines may be saving to a file, or even just generating internal state such as via `snprintf()`).

The problem of input formatting causing security problems is not an idle possibility; see CERT Advisory CA-2000-13 for an example of an exploit using this weakness. For more information on how these problems can be exploited, see Pascal Bouchareine's email article titled "[Paper] Format bugs", published in the July 18, 2000 edition of Bugtraq (<http://www.securityfocus.com>).

Of course, this all begs the question as to whether or not the internationalization lookup is, in fact, secure. If you're creating your own internationalization lookup routines, make sure that an untrusted user can only specify a legal locale and not something else like an arbitrary path.

Clearly, you want to limit the strings created through internationalization to ones you can trust. Otherwise, an attacker could use this ability to exploit the weaknesses in

format strings, particularly in C/C++ programs. This has been an item of discussion in Bugtraq (e.g., see John Levon's Bugtraq post on July 26, 2000). For more information, see the discussion on on permitting users to only select legal language values in Section 4.7.3

Although it's really a programming bug, it's worth mentioning that different countries notate numbers in different ways, in particular, both the period (.) and comma (,) are used to separate an integer from its fractional part. If you save or load data, you need to make sure that the active locale does not interfere with data handling. Otherwise, a French user may not be able to exchange data with an English user, because the data stored and retrieved will use different separators. I'm unaware of this being used as a security problem, but it's conceivable.

# Chapter 9. Language-Specific Issues

*Undoubtedly there are all sorts of languages in the world, yet none of them is without meaning.*  
*1 Corinthians 14:10 (NIV)*

There are many language-specific security issues. Many of them can be summarized as follows:

- Turn on all relevant warnings and protection mechanisms available to you where practical. For compiled languages, this includes both compile-time mechanisms and run-time mechanisms. In general, security-relevant programs should compile cleanly with all warnings turned on.
- Avoid dangerous and deprecated operations in the language. By “dangerous”, I mean operations which are difficult to use correctly.
- Ensure that the languages’ infrastructure (e.g., run-time library) is available and secured.
- Languages that automatically garbage-collect strings should be especially careful to immediately erase secret data (in particular secret keys and passwords).
- Know precisely the semantics of the operations that you are using. Look up operation’s semantics in its documentation. Do not ignore return values unless you’re sure they cannot be relevant. This is particularly difficult in languages which don’t support exceptions, like C, but that’s the way it goes.

## 9.1. C/C++

One of the biggest security problems with C and C++ programs is buffer overflow; see Chapter 5 for more information. C has the additional weakness of not supporting exceptions, which makes it easy to write programs that ignore critical error situations.

One complication in C and C++ is that the character type “char” can be signed or unsigned (depending on the compiler and machine). When a signed char with its high bit set is saved in an integer, the result will be a negative number; in some cases this can be exploitable. In general, use “unsigned char” instead of char or signed char for buffers, pointers, and casts when dealing with character data that may have values greater than 127 (0x7f).

C and C++ are by definition very lax in their type-checking support, but there’s no need to be lax in your code. Turn on as many compiler warnings as you can and change the code to cleanly compile with them, and strictly use ANSI prototypes in separate header (.h) files to ensure that all function calls use the correct types. For C or C++ compilations using gcc, use at least the following as compilation flags (which turn on a host of warning messages) and try to eliminate all warnings (note that -O2 is used since some warnings can only be detected by the data flow analysis performed at higher optimization levels):

```
gcc -Wall -Wpointer-arith -Wstrict-prototypes -O2
```

You might want “-W -pedantic” too.

Many C/C++ compilers can detect inaccurate format strings. For example, gcc can warn about inaccurate format strings for functions you create if you use its `__attribute__()` facility (a C extension) to mark such functions, and you can use that facility without making your code non-portable. Here is an example of what you’d put in your header (.h) file:

```
/* in header.h */
#ifdef __GNUC__
# define __attribute__(x) /*nothing*/
#endif

extern void logprintf(const char *format, ...)
    __attribute__((format(printf,1,2)));
extern void logprintva(const char *format, va_list args)
    __attribute__((format(printf,1,0)));
```

The “format” attribute takes either “printf” or “scanf”, and the numbers that follow are the parameter number of the format string and the first variadic parameter

(respectively). The GNU docs talk about this well. Note that there are other `__attribute__` facilities as well, such as "noreturn" and "const".

## 9.2. Perl

Perl programmers should first read the man page `perlsec(1)`, which describes a number of issues involved with writing secure programs in Perl. In particular, `perlsec(1)` describes the “taint” mode, which most secure Perl programs should use. Taint mode is automatically enabled if the real and effective user or group IDs differ, or you can use the `-T` command line flag (use the latter if you’re running on behalf of someone else, e.g., a CGI script). Taint mode turns on various checks, such as checking path directories to make sure they aren’t writable by others.

The most obvious affect of taint mode, however, is that you may not use data derived from outside your program to affect something else outside your program by accident. In taint mode, all externally-obtained input is marked as “tainted”, including command line arguments, environment variables, locale information (see `perllocale(1)`), results of certain system calls (`readdir`, `readlink`, the `gecos` field of `getpw*` calls), and all file input. Tainted data may not be used directly or indirectly in any command that invokes a sub-shell, nor in any command that modifies files, directories, or processes. There is one important exception: If you pass a list of arguments to either `system` or `exec`, the elements of that list are NOT checked for taintedness, so be especially careful with `system` or `exec` while in taint mode.

Any data value derived from tainted data becomes tainted also. There is one exception to this; the way to untaint data is to extract a substring of the tainted data. Don’t just use `“.*”` blindly as your substring, though, since this would defeat the tainting mechanism’s protections. Instead, identify patterns that identify the “safe” pattern allowed by your program, and use them to extract “good” values. After extracting the value, you may still need to check it (in particular for its length).

The `open`, `glob`, and backtick functions call the shell to expand filename wild card characters; this can be used to open security holes. You can try to avoid these functions entirely, or use them in a less-privileged “sandbox” as described in `perlsec(1)`. In

particular, backticks should be rewritten using the `system()` call (or even better, changed entirely to something safer).

The perl `open()` function comes with, frankly, “way too much magic” for most secure programs; it interprets text that, if not carefully filtered, can create lots of security problems. Before writing code to open or lock a file, consult the `perlopentut(1)` man page. In most cases, `sysopen()` provides a safer (though more convoluted) approach to opening a file. The new Perl 5.6 adds an `open()` call with 3 parameters to turn off the magic behavior without requiring the convolutions of `sysopen()` (<http://www.xray.mpe.mpg.de/mailling-lists/perl5-porters/2000-03/msg02596.html>).

Perl programs should turn on the warning flag (`-w`), which warns of potentially dangerous or obsolete statements.

You can also run Perl programs in a restricted environment. For more information see the “Safe” module in the standard Perl distribution. I’m uncertain of the amount of auditing that this has undergone, so beware of depending on this for security. You might also investigate the “Penguin Model for Secure Distributed Internet Scripting”, though at the time of this writing the code and documentation seems to be unavailable.

## 9.3. Python

As with any language, beware of any functions which allow data to be executed as parts of a program, to make sure an untrusted user can’t affect their input. This includes `exec()`, `eval()`, and `execfile()` (and frankly, you should check carefully any call to `compile()`). The `input()` statement is also surprisingly dangerous. [Watters 1996, 150].

Python programs with privileges that can be invoked by unprivileged users (e.g., `setuid/setgid` programs) must *not* import the “user” module. The user module causes the `pythonrc.py` file to be read and executed. Since this file would be under the control of an untrusted user, importing the user module allows an attacker to force the trusted program to run arbitrary code.

Python includes support for “Restricted Execution” through its `RExec` class. This is primarily intended for executing applets and mobile code, but it can also be used to

limit privilege in a program even when the code has not been provided externally. By default, a restricted execution environment permits reading (but not writing) of files, and does not include operations for network access or GUI interaction. These defaults can be changed, but beware of creating loopholes in the restricted environment. In particular, allowing a user to unrestrictedly add attributes to a class permits all sorts of ways to subvert the environment because Python's implementation calls many "hidden" methods. Note that, by default, most Python objects are passed by reference; if you insert a reference to a mutable value into a restricted program's environment, the restricted program can change the object in a way that's visible outside the restricted environment! Thus, if you want to give access to a mutable value, in many cases you should copy the mutable value or use the Bastion module (which supports restricted access to another object). For more information, see Kuchling [2000]. I'm uncertain of the amount of auditing that the restricted execution capability has undergone, so programmer beware.

## 9.4. Shell Scripting Languages (sh and csh Derivatives)

I strongly recommend against using standard command shell scripting languages (such as csh, sh, and bash) for setuid/setgid secure code. Some systems (such as Linux) completely disable them, so you're creating an unnecessary portability problem. On some old systems they are fundamentally insecure due to a race condition (as discussed in Section 3.1.3). Even for other systems, they're not really a good idea. Standard command shells are still notorious for being affected by nonobvious inputs - generally because command shells were designed to try to do things "automatically" for an interactive user, not to defend against a determined attacker. For example, "hidden" environment variables (e.g., the ENV or BASH\_ENV variable) can affect how they operate or even execute arbitrary user-defined code before the script can even execute. Even things like filenames of the executable or directory contents can affect things. For example, on many Bourne shell implementations, doing the following will grant root



access (thanks to NCSA for describing this exploit):

```
% ln -s /usr/bin/setuid-shell /tmp/-x
% cd /tmp
% -x
```

Some systems may have closed this hole, but the point still stands: most command shells aren't intended for writing secure setuid/setgid programs. For programming purposes, avoid creating setuid shell scripts, even on those systems that permit them. Instead, write a small program in another language to clean up the environment, then have it call other executables (some of which might be shell scripts).

If you still insist on using shell scripting languages, at least put the script in a directory where it cannot be moved or changed. Set PATH and IFS to known values very early in your script.

## 9.5. Ada

In Ada95, the `Unbounded_String` type is often more flexible than the `String` type because it is automatically resized as necessary. However, don't store especially sensitive values such as passwords or secret keys in an `Unbounded_String`, since core dumps and page areas might still hold them later. Instead, use the `String` type for this data and overwrite the data as soon as possible with some constant value such as others => ' '.

## 9.6. Java

If you're developing secure programs using Java, frankly your first step (after learning Java) is to read the two primary texts for Java security, namely Gong [1999] and McGraw [1999] (for the latter, look particularly at section 7.1). You should also look at Sun's posted security code guidelines at

<http://java.sun.com/security/seccodeguide.html>. A set of slides describing Java's security model are freely available at <http://www.dwheeler.com/javasec>.

Obviously, a great deal depends on the kind of application you're developing. Java code intended for use on the client side has a completely different environment (and trust model) than code on a server side. The general principles apply, of course; for example, you must check and filter any input from an untrusted source. However, in Java there are some "hidden" inputs or potential inputs that you need to be wary of, as discussed below. Johnathan Nightingale [2000] made an interesting statement summarizing many of the issues in Java programming:

... the big thing with Java programming is minding your inheritances. If you inherit methods from parents, interfaces, or parents' interfaces, you risk opening doors to your code.

The following are a few key guidelines, based on Gong [1999], McGraw [1999], Sun's guidance, and my own experience:

1. Do not use public fields or variables; declare them as private and provide accessors to them so you can limit their accessibility.
2. Make methods private unless there is a good reason to do otherwise (and if you do otherwise, document why). These non-private methods must protect themselves, because they may receive tainted data (unless you've somehow arranged to protect them).
3. The JVM may not actually enforce the accessibility modifiers (e.g., "private") at run-time in an application (as opposed to an applet). My thanks to John Steven (Cigital Inc.), who pointed this out on the "Secure Programming" mailing list on November 7, 2000. The issue is that it all depends on what class loader the class requesting the access was loaded with. If the class was loaded with a trusted class loader (including the null/ primordial class loader), the access check returns "TRUE" (allowing access). For example, this works (at least with Sun's 1.2.2 VM ; it might not work with other implementations):
  - a. write a victim class (V) with a public field, compile it.
  - b. write an 'attack' class (A) that accesses that field, compile it

- c. change V's public field to private, recompile
- d. run A - it'll access V's (now private) field.

However, the situation is different with applets. If you convert A to an applet and run it as an applet (e.g., with appletviewer or browser), its class loader is no longer a trusted (or null) class loader. Thus, the code will throw `java.lang.IllegalAccessError`, with the message that you're trying to access a field `V.secret` from class A.

4. Avoid using static field variables. Such variables are attached to the class (not class instances), and classes can be located by any other class. As a result, static field variables can be found by any other class, making them much more difficult to secure.
5. Never return a mutable object to potentially malicious code (since the code may decide to change it). Note that arrays are mutable (even if the array contents aren't), so don't return a reference to an internal array with sensitive data.
6. Never store user given mutable objects (including arrays of objects) directly. Otherwise, the user could hand the object to the secure code, let the secure code "check" the object, and change the data while the secure code was trying to use the data. Clone arrays before saving them internally, and be careful here (e.g., beware of user-written cloning routines).
7. Don't depend on initialization. There are several ways to allocate uninitialized objects.
8. Make everything final, unless there's a good reason not to. If a class or method is non-final, an attacker could try to extend it in a dangerous and unforeseen way. Note that this causes a loss of extensibility, in exchange for security.
9. Don't depend on package scope for security. A few classes, such as `java.lang`, are closed by default, and some Java Virtual Machines (JVMs) let you close off other packages. Otherwise, Java classes are not closed. Thus, an attacker could introduce a new class inside your package, and use this new class to access the things you thought you were protecting.

10. Don't use inner classes. When inner classes are translated into byte codes, the inner class is translated into a class accessible to any class in the package. Even worse, the enclosing class's private fields silently become non-private to permit access by the inner class!
11. Minimize privileges. Where possible, don't require any special permissions at all. McGraw goes further and recommends not signing any code; I say go ahead and sign the code (so users can decide to "run only signed code by this list of senders"), but try to write the program so that it needs nothing more than the sandbox set of privileges. If you must have more privileges, audit that code especially hard.
12. If you must sign your code, put it all in one archive file. Here it's best to quote McGraw [1999]:

The goal of this rule is to prevent an attacker from carrying out a mix-and-match attack in which the attacker constructs a new applet or library that links some of your signed classes together with malicious classes, or links together signed classes that you never meant to be used together. By signing a group of classes together, you make this attack more difficult. Existing code-signing systems do an inadequate job of preventing mix-and-match attacks, so this rule cannot prevent such attacks completely. But using a single archive can't hurt.

13. Make your classes uncloneable. Java's object-cloning mechanism allows an attacker to instantiate a class without running any of its constructors. To make your class uncloneable, just define the following method in each of your classes:

```
public final void clone() throws java.lang.CloneNotSupportedException {
    throw new java.lang.CloneNotSupportedException();
}
```

If you really need to make your class cloneable, then there are some protective measures you can take to prevent attackers from redefining your clone method. If you're defining your own clone method, just make it final. If you're not, you can at least prevent the clone method from being maliciously overridden by adding the following:

```
public final void clone() throws java.lang.CloneNotSupportedException {
    super.clone();
}
```

14. Make your classes unserializable. Serialization allows attackers to view the internal state of your objects, even private portions. To prevent this, add this method to your classes:

```
private final void writeObject(ObjectOutputStream out)
    throws java.io.IOException {
    throw new java.io.IOException("Object cannot be serialized");
}
```

Even in cases where serialization is okay, be sure to use the `transient` keyword for the fields that contain direct handles to system resources and that contain information relative to an address space. Otherwise, deserializing the class may permit improper access. You may also want to identify sensitive information as `transient`.

If you define your own serializing method for a class, it should not pass an internal array to any `DataInput/DataOutput` method that takes an array. The rationale: All `DataInput/DataOutput` methods can be overridden. If a `Serializable` class passes a private array directly to a `DataOutput(write(byte [] b))` method, then an attacker could subclass `ObjectOutputStream` and override the `write(byte [] b)` method to enable him to access and modify the private array. Note that the default serialization does not expose private byte array fields to `DataInput/DataOutput` byte array methods.

15. Make your classes underserializable. Even if your class is not serializable, it may still be deserializable. An attacker can create a sequence of bytes that happens to deserialize to an instance of your class with values of the attacker's choosing. In other words, deserialization is a kind of public constructor, allowing an attacker to choose the object's state - clearly a dangerous operation! To prevent this, add this method to your classes:

```
private final void readObject(ObjectInputStream in)
    throws java.io.IOException {
    throw new java.io.IOException("Class cannot be deserialized");
}
```

16. Don't compare classes by name. After all, attackers can define classes with identical names, and if you're not careful you can cause confusion by granting

these classes undesirable privileges. Thus, here's an example of the *wrong* way to determine if an object has a given class:

```
if (obj.getClass().getName().equals("Foo")) {
```

If you need to determine if two objects have exactly the same class, instead use `getClass()` on both sides and compare using the `==` operator. Thus, you should use this form:

```
if (a.getClass() == b.getClass()) {
```

If you truly need to determine if an object has a given classname, you need to be pedantic and be sure to use the current namespace (of the current class's `ClassLoader`). Thus, you'll need to use this format:

```
if (obj.getClass() == this.getClassLoader().loadClass("Foo")) {
```

This guideline is from McGraw and Felten, and it's a good guideline. I'll add that, where possible, it's often a good idea to avoid comparing class values anyway. It's often better to try to design class methods and interfaces so you don't need to do this at all. However, this isn't always practical, so it's important to know these tricks.

17. Don't store secrets (cryptographic keys, passwords, or algorithm) in the code or data. Hostile JVMs can quickly view this data. Code obfuscation doesn't really hide the code from serious attackers.

## 9.7. TCL

Tcl stands for "tool command language" and is pronounced "tickle." TCL is divided into two parts: a language and a library. The language is a simple text language, intended for issuing commands to interactive programs and including basic programming capabilities. The library can be embedded in application programs.

You can find more information about TCL at sites such as the TCL WWW Info (<http://www.sco.com/Technology/tcl/Tcl.html>) web page. Probably of most interest are

Safe-TCL (which creates a sandbox in TCL) and Safe-TK (which implements a sandboxed portable GUI for Safe-TCL), as well as the WebWiseTclTk Toolkit permits TCL packages to be automatically located and loaded from anywhere on the World Wide Web. You can find more about the latter from <http://www.cbl.ncsu.edu/software/WebWiseTclTk>. It's not clear to me how much code review this has received. More useful information is available from the comp.lang.tcl FAQ launch page at <http://www.tclfaq.wservice.com/tcl-faq>. However, it's worth noting that TCL's desire to be a small, "simple" language results in a language that can be rather limiting; see Richard Stallman's "Why You Should Not Use TCL" (<http://sdg.lcs.mit.edu/~jchapin/6853-FT97/Papers/stallman-tcl.html>). For example, TCL's notion that there is essentially only one data type (string) can make many programs harder to write (as well as making them slow). Also, when I've written TCL programs I've found that it's easy to accidentally create TCL programs where malicious input strings can cause untoward and unexpected behavior. For example, an attackers may be able to cause your TCL program to do unexpected things by sending characters with special meaning to TCL such as embedded spaces, double-quote, curly braces, dollar signs, or brackets (or create input to cause these characters to be created during processing). Thus, I don't recommend TCL for writing programs which must mediate a security boundary. If you do choose to do so, be especially careful to ensure that user input cannot "fool" the program. On the other hand, I know of no strong reason (other than insufficient review) that TCL programs can't be used to implement mobile code. There are certainly TCL advocates who will advocate more use than I do, and TCL is one of the few languages with a ready-made sandbox implementation.

# Chapter 10. Special Topics

*Understanding is a fountain of  
life to those who have it, but folly  
brings punishment to fools.  
Proverbs 16:22 (NIV)*

## 10.1. Passwords

Where possible, don't write code to handle passwords. In particular, if the application is local, try to depend on the normal login authentication by a user. If the application is a CGI script, try to depend on the web server to provide the protection. If the application is over a network, avoid sending the password as cleartext (where possible) since it can be easily captured by network sniffers and reused later. "Encrypting" a password using some key fixed in the algorithm or using some sort of shrouding algorithm is essentially the same as sending the password as cleartext.

For networks, consider at least using digest passwords. Digest passwords are passwords developed from hashes; typically the server will send the client some data (e.g., date, time, name of server), the client combines this data with the user password, the client hashes this value (termed the "digest password") and replies just the hashed result to the server; the server verifies this hash value. This works, because the password is never actually sent in any form; the password is just used to derive the hash value. Digest passwords aren't considered "encryption" in the usual sense and are usually accepted even in countries with laws constraining encryption for confidentiality. Digest passwords are vulnerable to active attack threats but protect against passive network sniffers. One weakness is that, for digest passwords to work, the server must have all the unhashed passwords, making the server a very tempting target for attack.

If your application permits users to set their passwords, check the passwords and permit only "good" passwords (e.g., not in a dictionary, having certain minimal length, etc.). You may want to look at information such as



[http://consult.cern.ch/writeup/security/security\\_3.html](http://consult.cern.ch/writeup/security/security_3.html) on how to choose a good password. You should use PAM if you can, because it supports pluggable password checkers.

## 10.2. Random Numbers

“Random” numbers generated by many library routines are intended to be used for simulations, games, and so on; they are *not* sufficiently random for use in security functions such as key generation. The problem is that these library routines use algorithms whose future values can be easily deduced by an attacker (though they may appear random). For security functions, you need random values based on truly unpredictable values such as quantum effects.

Failing to correctly generate truly random values for keys has caused a number of problems, including holes in Kerberos, the X window system, and NFS [Venema 1996].

The Linux kernel (since 1.3.30) includes a random number generator, which is sufficient for many security purposes. This random number generator gathers environmental noise from device drivers and other sources into an entropy pool. When accessed as `/dev/random`, random bytes are only returned within the estimated number of bits of noise in the entropy pool (when the entropy pool is empty, the call blocks until additional environmental noise is gathered). When accessed as `/dev/urandom`, as many bytes as are requested are returned even when the entropy pool is exhausted. If you are using the random values for cryptographic purposes (e.g., to generate a key), use `/dev/random`. More information is available in the system documentation `random(4)`.

## 10.3. Specially Protect Secrets (Passwords and Keys) in User Memory

If your application must handle passwords or non-public keys (such as session keys, private keys, or secret keys), overwrite them immediately after using them so they have

minimal exposure. For example, in Java, don't use the type `String` to store a password because `Strings` are immutable (they will not be overwritten until garbage-collected and reused, possibly a far time in the future). Instead, in Java use `char[]` to store a password, so it can be immediately overwritten.

Also, if your program handles such secret values, be sure to disable creating core dumps (via `ulimit`). Otherwise, an attacker may be able to halt the program and find the secret value in the data dump. Also, beware - normally processes can monitor other processes through the calls for debuggers (e.g., via `ptrace(2)` and the `/proc` pseudo-filesystem) [Venema 1996] Kernels usually protect against these monitoring routines if the process is `setuid` or `setgid` (on the few ancient ones that don't, there really isn't a way to defend yourself other than upgrading). Thus, if your process manages secret values, you probably should make it `setgid` or `setuid` (to a different unprivileged group or user) to forceably inhibit this kind of monitoring.

## 10.4. Cryptographic Algorithms and Protocols

Often cryptographic algorithms and protocols are necessary to keep a system secure, particularly when communicating through an untrusted network such as the Internet. Where possible, use session encryption to foil session hijacking and to hide authentication information, as well as to support privacy.

For background information and code, you should probably look at the classic text "Applied Cryptography" [Schneier 1996]. The newsgroup "sci.crypt" has a series of FAQ's; you can find them at many locations, including <http://www.landfield.com/faqs/cryptography-faq>. Linux-specific resources include the Linux Encryption HOWTO at <http://marc.mutz.com/Encryption-HOWTO/>. A discussion on how protocols use the basic algorithms can be found in [Opplinger 1998]. What follows here is just a few comments; these areas are rather specialized and covered more thoroughly elsewhere.

It's worth noting that there are many legal hurdles involved with cryptographic algorithms. First, the use, export, and/or import of implementations of encryption algorithms are restricted in many countries. Second, a number of algorithms are

patented; even if the owners permit “free use” at the moment, without a signed contract they can always change their minds later. Most of the patent issues can be easily avoided nowadays, once you know to watch out for it, so there’s little reason to subject yourself to the problem.

Cryptographic protocols and algorithms are difficult to get right, so do not create your own. Instead, use existing protocols and algorithms where you can. In particular, do not create your own encryption algorithms unless you are an expert in cryptology, know what you’re doing, and plan to spend years in professional review of the algorithm. Creating encryption algorithms (that are any good) is a task for experts only.

For protocols, try to use standard-conforming protocols such as SSL (soon to be TLS), SSH, IPSec, GnuPG/PGP, and Kerberos. Many of these overlap somewhat in functionality, but each has a “specialty” niche. SSL (soon to be TLS) is the primary method for protecting http (web) transactions. PGP-compatible protocols (implemented in PGP and GnuPG) are a primary method for securing email end-to-end. Kerberos is a primary method for securing and supporting authentication on a LAN. SSH is the primary method of securing “remote terminals” over an internet, e.g., telnet-like and X windows connections, though it’s often used for securing other data streams too (such as CVS accesses). IPSec is the primary method for security lower-level packets and “all” packets, so it’s particularly useful for securing virtual private networks and remote machines.

For secret key (bulk data) encryption algorithms, use only encryption algorithms that have been openly published and withstood years of attack, and check on their patent status. For encrypting unimportant data, the old DES (56-bit key) algorithm still has some value, but with modern hardware it’s too easy to break. For many applications triple-DES is currently the best encryption algorithm; it has a reasonably lengthy key (112 bits), no patent issues, and has a long history of withstanding attacks. The upcoming AES algorithm may be worth using as well, once it’s proven. You should probably avoid IDEA due to patent issues (it’s subject to U.S. and European patents), but I’m unaware of any serious technical problems with it. Your protocol should support multiple algorithms; that way, when an algorithm is broken, users can switch to another one.

For public key cryptography (used, among other things, for authentication and sending

secret keys), there are only a few widely-deployed algorithms. One of the most widely-used algorithms is RSA; RSA's algorithm is patented, but only in the U.S., and that patent expires September 20, 2000. The Diffie-Hellman key exchange algorithm is widely used to permit two parties to agree on a session key. By itself it doesn't guarantee that the parties are who they say they are, or that there is no middleman, but it does strongly help defend against passive listeners; its patent expired in 1997. NIST developed the digital signature standard (DSS) (it's a modification of the ElGamal cryptosystem) for digital signature generation and verification; one of the conditions for its development was for it to be patent-free.

Some programs need a one-way hash algorithm, that is, a function that takes an "arbitrary" amount of data and generates a fixed-length number that hard to invert (e.g., it's difficult for an attacker to create a different set of data to generate that same value). For a number of years MD5 has been a favorite, but recent efforts have shown that its 128-bit length may not be enough [van Oorschot 1994] and that certain attacks weaken MD5' protection [Dobbertin 1996]. Indeed, there are rumors that a top industry cryptographer has broken MD5, but is bound by employee agreement to keep silent (see the Bugtraq 22 August 2000 posting by John Viega). If you're writing new code, you probably ought to use SHA-1 instead.

In a related note, if you must create your own communication protocol, examine the problems of what's gone on before. Classics such as Bellovin [1989]'s review of security problems in the TCP/IP protocol suite might help you, as well as Bruce Schneier [1998] and Mudge's breaking of Microsoft's PPTP implementation and their follow-on work. Of course, be sure to give any new protocol widespread review, and reuse what you can.

## 10.5. PAM

Pluggable Authentication Modules (PAM) is a flexible mechanism for authenticating users. Many Unix-like systems support PAM, including Solaris, nearly all Linux distributions (e.g., Red Hat Linux, Caldera, and Debian as of version 2.2), and FreeBSD as of version 3.1. By using PAM, your program can be independent of the

authentication scheme (passwords, SmartCards, etc.). Basically, your program calls PAM, which at run-time determines which “authentication modules” are required by checking the configuration set by the local system administrator. If you’re writing a program that requires authentication (e.g., entering a password), you should include support for PAM. You can find out more about the Linux-PAM project at <http://www.kernel.org/pub/linux/libs/pam/index.html>.

## 10.6. Tools

Some tools may help you detect security problems before you field the result. If you’re building a common kind of product where many standard potential flaws exist (like an ftp server or firewall), you might find standard security scanning tools useful. One good one is Nessus (<http://www.nessus.org>); there are many others. Of course, running a “secure” program on an insecure platform configuration makes little sense; you may want to examine hardening systems such as Bastille available at <http://www.bastille-linux.org>.

You may find some auditing tools helpful for finding potential security flaws. Here are a few:

- ITS4 from Cigital (formerly Reliable Software Technologies, RST) statically checks C/C++ code. ITS4 works by performing pattern-matching on source code, looking for patterns known to be possibly dangerous (e.g., certain function calls). It is available free for non-commercial use, including its source code and with certain modification and redistribution rights. One warning; the tool’s licensing claims can be initially misleading. Cigital claims that ITS4 is “open source” but, in fact, its license does not meet the Open Source Definition (<http://www.opensource.org/osd.html>) (OSD). In particular, ITS4’s license fails point 6, which forbids “non-commercial use only” clauses in open source licenses. It’s unfortunate that Cigital insists on using the term “open source” to describe their license. ITS4 is a fine tool, released under a fairly generous license for commercial software, yet using the term this way can give the appearance of a company trying to gain the cachet of “open source” without actually being open source. Cigital says

that they simply don't accept the OSD definition and that they wish to use a different definition instead. Nothing legally prevents this, but the OSD definition is used by over 5000 software projects (at least all those hosted by SourceForge at <http://www.sourceforge.net>), Linux distributors, Netscape (now AOL), the W3C, journalists (such as those of the Economist), and many other organizations. Most programmers don't want to wade through license agreements, so using this other definition can be confusing. I do not believe Cigital has any intention to mislead; they're a reputable company with very reputable and honest people. It's unfortunate that this particular position of theirs leads (in my opinion) to unnecessary confusion. In any case, ITS4 is available at <http://www.rstcorp.com/its4>.

- LCLint is a tool for statically checking C programs. With minimal effort, LCLint can be used as a better lint. If additional effort is invested adding annotations to programs, LCLint can perform stronger checking than can be done by any standard lint. The software is licensed under the GPL and is available from <http://lclint.cs.virginia.edu>.
- BFBTester, the Brute Force Binary Tester, is licensed under the GPL. This program does quick security checks of binary programs. BFBTester performs checks of single and multiple argument command line overflows and environment variable overflows. Version 2.0 and higher can also watch for tempfile creation activity (to check for using unsafe tempfile names). More information is available at <http://my.ispchannel.com/~mheffner/bfbtester>.

## 10.7. Miscellaneous

The following are miscellaneous security guidelines that I couldn't seem to fit anywhere else:

Have your program check at least some of its assumptions before it uses them (e.g., at the beginning of the program). For example, if you depend on the "sticky" bit being set on a given directory, test it; such tests take little time and could prevent a serious problem. If you worry about the execution time of some tests on each call, at least perform the test at installation time, or even better at least perform the test on

application start-up.

Write audit logs for program startup, session startup, and for suspicious activity. Possible information of value includes date, time, uid, euid, gid, egid, terminal information, process id, and command line values. You may find the function `syslog(3)` helpful for implementing audit logs. One awkward problem is that any logging system should be able to record a lot of information (since this information could be very helpful), yet if the information isn't handled carefully the information itself could be used to create an attack. After all, the attacker controls some of the input being sent to the program. When recording data sent by a possible attacker, identify a list of "expected" characters and escape any "unexpected" characters so that the log isn't corrupted. Not doing this can be a real problem; users may include characters such as control characters (especially `NIL` or end-of-line) that can cause real problems. For example, if an attacker embeds a newline, they can then forge log entries by following the newline with the desired log entry. Sadly, there doesn't seem to be standard convention for escaping these characters. I'm partial to the URL escaping mechanism (`%hh` where `hh` is the hexadecimal value of the escaped byte) but there are others including the C convention (`\ooo` for the octal value and `\X` where `X` is a special symbol, e.g., `\n` for newline). There's also the caret-system (`^I` is control-I), though that doesn't handle byte values over 127 gracefully.

There is the danger that a user could create a denial-of-service attack (or at least stop auditing) by performing a very large number of events that cut an audit record until the system runs out of resources to store the records. One approach to counter to this threat is to rate-limit audit record recording; intentionally slow down the response rate if "too many" audit records are being cut. You could try to slow the response rate only to the suspected attacker, but in many situations a single attacker can masquerade as potentially many users.

Selecting what is "suspicious activity" is, of course, dependent on what the program does and its anticipated use. Any input that fails the filtering checks discussed earlier is certainly a candidate (e.g., containing `NIL`). Inputs that could not result from normal use should probably be logged, e.g., a CGI program where certain required fields are missing in suspicious ways. Any input with phrases like `/etc/passwd` or `/etc/shadow` or the like is very suspicious in many cases. Similarly, trying to access Windows "registry" files or `.pwl` files is very suspicious.

If you have a built-in scripting language, it may be possible for the language to set an environment variable which adversely affects the program invoking the script. Defend against this.

If you need a complex configuration language, make sure the language has a comment character and include a number of commented-out secure examples. Often '#' is used for commenting, meaning "the rest of this line is a comment".

If possible, don't create setuid or setgid root programs; make the user log in as root instead.

Sign your code. That way, others can check to see if what's available was what was sent.

Consider statically linking secure programs. This counters attacks on the dynamic link library mechanism by making sure that the secure programs don't use it.

When reading over code, consider all the cases where a match is not made. For example, if there is a switch statement, what happens when none of the cases match? If there is an "if" statement, what happens when the condition is false?



# Chapter 11. Conclusion

*The end of a matter is better than its beginning, and patience is better than pride.*

*Ecclesiastes 7:8 (NIV)*

Designing and implementing a truly secure program is actually a difficult task on Unix-like systems such as Linux and Unix. The difficulty is that a truly secure program must respond appropriately to all possible inputs and environments controlled by a potentially hostile user. Developers of secure programs must deeply understand their platform, seek and use guidelines (such as these), and then use assurance processes (such as peer review) to reduce their programs' vulnerabilities.

In conclusion, here are some of the key guidelines from this paper:

- Validate all your inputs, including command line inputs, environment variables, CGI inputs, and so on. Don't just reject "bad" input; define what is an "acceptable" input and reject anything that doesn't match.
- Avoid buffer overflow. Make sure that long inputs (and long intermediate data values) can't be used to take over your program. This is the primary programmatic error at this time.
- Structure Program Internals. Secure the interface, minimize privileges, make the initial configuration and defaults safe, and fail safe. Avoid race conditions (e.g., by safely opening any files in a shared directory like /tmp). Trust only trustworthy channels (e.g., most servers must not trust their clients for security checks or other sensitive data such as an item's price in a purchase).
- Carefully call out to other resources. Limit their values to valid values (in particular be concerned about metacharacters), and check all system call return values.
- Reply information judiciously. In particular, minimize feedback, and handle full or unresponsive output to an untrusted user.

# Chapter 12. Bibliography

*The words of the wise are like  
goads, their collected sayings like  
firmly embedded nails—given by  
one Shepherd. Be warned, my  
son, of anything in addition to  
them. Of making many books  
there is no end, and much study  
wearies the body.*

*Ecclesiastes 12:11-12 (NIV)*

*Note that there is a heavy emphasis on technical articles available on the web, since this is where most of this kind of technical information is available.*

[Advosys 2000] Advosys Consulting (formerly named Webber Technical Services). *Writing Secure Web Applications*. <http://advosys.ca/tips/web-security.html>

[Al-Herbish 1999] Al-Herbish, Thamer. 1999. *Secure Unix Programming FAQ*. <http://www.whitefang.com/sup>.

[Aleph1 1996] Aleph1. November 8, 1996. “Smashing The Stack For Fun And Profit”. *Phrack Magazine*. Issue 49, Article 14. <http://www.phrack.com/search.phtml?view&article=p49-14> or alternatively <http://www.2600.net/phrack/p49-14.html>.

[Anonymous 1999] Anonymous. October 1999. *Maximum Linux Security: A Hacker’s Guide to Protecting Your Linux Server and Workstation* Sams. ISBN: 0672316706.

[Anonymous 1998] Anonymous. September 1998. *Maximum Security : A Hacker’s Guide to Protecting Your Internet Site and Network*. Sams. Second Edition. ISBN: 0672313413.

[AUSCERT 1996] Australian Computer Emergency Response Team (AUSCERT) and O’Reilly. May 23, 1996 (rev 3C). *A Lab Engineers Check List for Writing Secure Unix Code*. [ftp://ftp.auscert.org.au/pub/auscert/papers/secure\\_programming\\_checklist](ftp://ftp.auscert.org.au/pub/auscert/papers/secure_programming_checklist)

- [Bach 1986] Bach, Maurice J. 1986. *The Design of the Unix Operating System*. Englewood Cliffs, NJ: Prentice-Hall, Inc. ISBN 0-13-201799-7 025.
- [Bellovin 1989] Bellovin, Steven M. April 1989. "Security Problems in the TCP/IP Protocol Suite" *Computer Communications Review* 2:19, pp. 32-48.  
<http://www.research.att.com/~smb/papers/ipext.pdf>
- [Bellovin 1994] Bellovin, Steven M. December 1994. *Shifting the Odds – Writing (More) Secure Software*. Murray Hill, NJ: AT&T Research.  
<http://www.research.att.com/~smb/talks>
- [Bishop 1996] Bishop, Matt. May 1996. "UNIX Security: Security in Programming". *SANS '96*. Washington DC (May 1996).  
<http://olympus.cs.ucdavis.edu/~bishop/secprog.html>
- [Bishop 1997] Bishop, Matt. October 1997. "Writing Safe Privileged Programs". *Network Security 1997* New Orleans, LA.  
<http://olympus.cs.ucdavis.edu/~bishop/secprog.html>
- [CC 1999] *The Common Criteria for Information Technology Security Evaluation (CC)*. August 1999. Version 2.1. Technically identical to International Standard ISO/IEC 15408:1999. <http://csrc.nist.gov/cc/ccv20/ccv2list.htm>
- [CERT 1998] Computer Emergency Response Team (CERT) Coordination Center (CERT/CC). February 13, 1998. *Sanitizing User-Supplied Data in CGI Scripts*. CERT Advisory CA-97.25.CGI\_metachar.  
[http://www.cert.org/advisories/CA-97.25.CGI\\_metachar.html](http://www.cert.org/advisories/CA-97.25.CGI_metachar.html).
- [CMU 1998] Carnegie Mellon University (CMU). February 13, 1998 Version 1.4. "How To Remove Meta-characters From User-Supplied Data In CGI Scripts".  
[ftp://ftp.cert.org/pub/tech\\_tips/cgi\\_metacharacters](ftp://ftp.cert.org/pub/tech_tips/cgi_metacharacters).
- [Cowan 1999] Cowan, Crispin, Perry Wagle, Calton Pu, Steve Beattie, and Jonathan Walpole. "Buffer Overflows: Attacks and Defenses for the Vulnerability of the Decade". Proceedings of DARPA Information Survivability Conference and Expo (DISCEX), <http://schafercorp-ballston.com/discex> To appear at SANS 2000, <http://www.sans.org/newlook/events/sans2000.htm>. For a copy, see <http://immunix.org/documentation.html>.

- [Dobbertin 1996]. Dobbertin, H. 1996. The Status of MD5 After a Recent Attack. RSA Laboratories' CryptoBytes. Vol. 2, No. 2.
- [Fenzi 1999] Fenzi, Kevin, and Dave Wrenski. April 25, 1999. *Linux Security HOWTO*. Version 1.0.2. <http://www.linuxdoc.org/HOWTO/Security-HOWTO.html>
- [FHS 1997] Filesystem Hierarchy Standard (FHS 2.0). October 26, 1997. Filesystem Hierarchy Standard Group, edited by Daniel Quinlan. Version 2.0. <http://www.pathname.com/fhs>.
- [Filipski 1986] Filipski, Alan and James Hanko. April 1986. "Making Unix Secure." Byte (Magazine). Peterborough, NH: McGraw-Hill Inc. Vol. 11, No. 4. ISSN 0360-5280. pp. 113-128.
- [FOLDOC] Free On-Line Dictionary of Computing. <http://foldoc.doc.ic.ac.uk/foldoc/index.html>.
- [FreeBSD 1999] FreeBSD, Inc. 1999. "Secure Programming Guidelines". *FreeBSD Security Information*. <http://www.freebsd.org/security/security.html>
- [FSF 1998] Free Software Foundation. December 17, 1999. *Overview of the GNU Project*. <http://www.gnu.ai.mit.edu/gnu/gnu-history.html>
- [FSF 1999] Free Software Foundation. January 11, 1999. *The GNU C Library Reference Manual*. Edition 0.08 DRAFT, for Version 2.1 Beta of the GNU C Library. Available at, for example, [http://www.netppl.fi/~pp/glibc21/libc\\_toc.html](http://www.netppl.fi/~pp/glibc21/libc_toc.html)
- [Galvin 1998a] Galvin, Peter. April 1998. "Designing Secure Software". *Sunworld*. <http://www.sunworld.com/swol-04-1998/swol-04-security.html>.
- [Galvin 1998b] Galvin, Peter. August 1998. "The Unix Secure Programming FAQ". *Sunworld*. <http://www.sunworld.com/sunworldonline/swol-08-1998/swol-08-security.html>
- [Garfinkel 1996] Garfinkel, Simson and Gene Spafford. April 1996. *Practical UNIX & Internet Security, 2nd Edition*. ISBN 1-56592-148-8. Sebastopol, CA: O'Reilly & Associates, Inc. <http://www.oreilly.com/catalog/puis>
- [Garfinkle 1997] Garfinkle, Simson. August 8, 1997. 21 Rules for Writing Secure CGI Programs. <http://webreview.com/wr/pub/97/08/08/bookshelf>

- [Graham 1999] Graham, Jeff. May 4, 1999. *Security-Audit's Frequently Asked Questions (FAQ)*. <http://lsap.org/faq.txt>
- [Gong 1999] Gong, Li. June 1999. *Inside Java 2 Platform Security*. Reading, MA: Addison Wesley Longman, Inc. ISBN 0-201-31000-7.
- [Gundavaram Unknown] Gundavaram, Shishir, and Tom Christiansen. Date Unknown. *Perl CGI Programming FAQ*.  
<http://language.perl.com/CPAN/doc/FAQs/cgi/perl-cgi-faq.html>
- [Hall "Beej" 1999] Hall, Brian "Beej". Beej's Guide to Network Programming Using Internet Sockets. 13-Jan-1999. Version 1.5.5.  
<http://www.ecst.csuchico.edu/~beej/guide/net>
- [Jones 2000] Jones, Jennifer. October 30, 2000. "Banking on Privacy". InfoWorld, Volume 22, Issue 44. San Mateo, CA: International Data Group (IDG). pp. 1-12.
- [Kernighan 1988] Kernighan, Brian W., and Dennis M. Ritchie. 1988. *The C Programming Language*. Second Edition. Englewood Cliffs, NJ: Prentice-Hall. ISBN 0-13-110362-8.
- [Kim 1996] Kim, Eugene Eric. 1996. *CGI Developer's Guide*. SAMS.net Publishing. ISBN: 1-57521-087-8 <http://www.eekim.com/pubs/cgibook>
- Kuchling [2000]. Kuchling, A.M. 2000. Restricted Execution HOWTO.  
<http://www.python.org/doc/howto/rexec/rexec.html>
- [McClure 1999] McClure, Stuart, Joel Scambray, and George Kurtz. 1999. *Hacking Exposed: Network Security Secrets and Solutions*. Berkeley, CA: Osbourne/McGraw-Hill. ISBN 0-07-212127-0.
- [McKusick 1999] McKusick, Marshall Kirk. January 1999. "Twenty Years of Berkeley Unix: From AT&T-Owned to Freely Redistributable." *Open Sources: Voices from the Open Source Revolution*.  
<http://www.oreilly.com/catalog/opensources/book/kirkmck.html>.
- [McGraw 1999] McGraw, Gary, and Edward W. Felten. January 25, 1999. *Securing Java: Getting Down to Business with Mobile Code*, 2nd Edition John Wiley & Sons. ISBN 047131952X. <http://www.securingsjava.com>.

- [McGraw 2000] McGraw, Gary and John Viega. March 1, 2000. Make Your Software Behave: Learning the Basics of Buffer Overflows.  
<http://www-4.ibm.com/software/developer/library/overflows/index.html>.
- [Miller 1995] Miller, Barton P., David Koski, Cjin Pheow Lee, Vivekananda Maganty, Ravi Murthy, Ajitkumar Natarajan, and Jeff Steidl. 1995. Fuzz Revisited: A Re-examination of the Reliability of UNIX Utilities and Services.  
[ftp://grilled.cs.wisc.edu/technical\\_papers/fuzz-revisited.pdf](ftp://grilled.cs.wisc.edu/technical_papers/fuzz-revisited.pdf).
- [Miller 1999] Miller, Todd C. and Theo de Raadt. “strncpy and strlcat – Consistent, Safe, String Copy and Concatenation” *Proceedings of Usenix '99*.  
<http://www.usenix.org/events/usenix99/millert.html> and  
[http://www.usenix.org/events/usenix99/full\\_papers/millert/PACKING\\_LIST](http://www.usenix.org/events/usenix99/full_papers/millert/PACKING_LIST)
- [Mudge 1995] Mudge. October 20, 1995. *How to write Buffer Overflows*. 10pht advisories. <http://www.10pht.com/advisories/bufero.html>.
- [NCSA] NCSA Secure Programming Guidelines.  
<http://www.ncsa.uiuc.edu/General/Grid/ACES/security/programming>.
- [Open Group 1997] The Open Group. 1997. *Single UNIX Specification, Version 2 (UNIX 98)*. <http://www.opengroup.org/online-pubs?DOC=007908799>.
- [OSI 1999]. Open Source Initiative. 1999. *The Open Source Definition*.  
<http://www.opensource.org/osd.html>.
- [Opplinger 1998] Opplinger, Rolf. 1998. Internet and Intranet Security. Norwood, MA: Artech House. ISBN 0-89006-829-1.
- [Peteanu 2000] Peteanu, Razvan. July 18, 2000. Best Practices for Secure Web Development. <http://members.home.net/razvan.peteanu>
- [Pfleeger 1997] Pfleeger, Charles P. 1997. *Security in Computing*. Upper Saddle River, NJ: Prentice-Hall PTR. ISBN 0-13-337486-6.
- [Phillips 1995] Phillips, Paul. September 3, 1995. *Safe CGI Programming*.  
<http://www.go2net.com/people/paulp/cgi-security/safe-cgi.txt>
- [Quintero 1999] Quintero, Federico Mena, Miguel de Icaza, and Morten Welinder  
GNOME Programming Guidelines

<http://developer.gnome.org/doc/guides/programming-guidelines/book1.html>

[Raymond 1997] Raymond, Eric. 1997. *The Cathedral and the Bazaar*.

<http://www.tuxedo.org/~esr/writings/cathedral-bazaar>

[Raymond 1998] Raymond, Eric. April 1998. *Homesteading the Noosphere*.

<http://www.tuxedo.org/~esr/writings/homesteading/homesteading.html>

[Ranum 1998] Ranum, Marcus J. 1998. *Security-critical coding for programmers - a C and UNIX-centric full-day tutorial*. <http://www.clark.net/pub/mjr/pubs/pdf/>.

[RFC 822] August 13, 1982 *Standard for the Format of ARPA Internet Text Messages*.

IETF RFC 822. <http://www.ietf.org/rfc/rfc0822.txt>.

[rfp 1999]. rain.forest.puppy. “Perl CGI problems”. *Phrack Magazine*. Issue 55, Article 07. <http://www.phrack.com/search.phtml?view&article=p55-7> or

<http://www.insecure.org/news/P55-07.txt>.

[Rochkind 1985]. Rochkind, Marc J. *Advanced Unix Programming*. Englewood Cliffs, NJ: Prentice-Hall, Inc. ISBN 0-13-011818-4.

[St. Laurent 2000] St. Laurent, Simon. February 2000. *XTech 2000 Conference Reports*.

“When XML Gets Ugly”. <http://www.xml.com/pub/2000/02/xtech/megginson.html>.

[Saltzer 1974] Saltzer, J. July 1974. “Protection and the Control of Information Sharing in MULTICS”. *Communications of the ACM*. v17 n7. pp. 388-402.

[Saltzer 1975] Saltzer, J., and M. Schroeder. September 1975. “The Protection of Information in Computing Systems”. *Proceedings of the IEEE*. v63 n9. pp. 1278-1308.

<http://www.mediacity.com/~norm/CapTheory/ProtInf>. Summarized in [Pfleeger 1997, 286].

[Schneier 1996] Schneier, Bruce. 1996. *Applied Cryptography, Second Edition: Protocols, Algorithms, and Source Code in C*. New York: John Wiley and Sons. ISBN 0-471-12845-7.

[Schneier 1998] Schneier, Bruce and Mudge. November 1998. *Cryptanalysis of Microsoft's Point-to-Point Tunneling Protocol (PPTP)* Proceedings of the 5th ACM Conference on Communications and Computer Security, ACM Press.

<http://www.counterpane.com/pptp.html>.

- [Schneier 1999] Schneier, Bruce. September 15, 1999. "Open Source and Security". *Crypto-Gram*. Counterpane Internet Security, Inc.  
<http://www.counterpane.com/crypto-gram-9909.html>
- [Seifried 1999] Seifried, Kurt. October 9, 1999. *Linux Administrator's Security Guide*.  
<http://www.securityportal.com/lasg>.
- [Shankland 2000] Shankland, Stephen. "Linux poses increasing threat to Windows 2000". CNET. <http://news.cnet.com/news/0-1003-200-1549312.html>
- [Shostack 1999] Shostack, Adam. June 1, 1999. *Security Code Review Guidelines*.  
<http://www.homeport.org/~adam/review.html>.
- [Sibert 1996] Sibert, W. Olin. Malicious Data and Computer Security. (NIST) NISSC '96. <http://www.fish.com/security/maldata.html>
- [Sitaker 1999] Sitaker, Kragen. Feb 26, 1999. *How to Find Security Holes*  
<http://www.pobox.com/~kragen/security-holes.html> and  
<http://www.dnaco.net/~kragen/security-holes.html>
- [SSE-CMM 1999] SSE-CMM Project. April 1999. *System Security Engineering Capability Maturity Model (SSE CMM) Model Description Document*. Version 2.0.  
<http://www.sse-cmm.org>
- [Stein 1999]. Stein, Lincoln D. September 13, 1999. *The World Wide Web Security FAQ*. Version 2.0.1 <http://www.w3.org/Security/Faq/www-security-faq.html>
- [Thompson 1974] Thompson, K. and D.M. Richie. July 1974. "The UNIX Time-Sharing System". *Communications of the ACM* Vol. 17, No. 7. pp. 365-375.
- [Torvalds 1999] Torvalds, Linus. February 1999. "The Story of the Linux Kernel". *Open Sources: Voices from the Open Source Revolution*. Edited by Chris Dibona, Mark Stone, and Sam Ockman. O'Reilly and Associates. ISBN 1565925823.  
<http://www.oreilly.com/catalog/opensources/book/linus.html>
- [Unknown] *SETUID(7)* <http://www.homeport.org/~adam/setuid.7.html>.
- [Van Biesbrouck 1996] Van Biesbrouck, Michael. April 19, 1996.  
<http://www.csclub.uwaterloo.ca/u/mlvanbie/cgisec>.



- [van Oorschot 1994] van Oorschot, P. and M. Wiener. November 1994. "Parallel Collision Search with Applications to Hash Functions and Discrete Logarithms." Proceedings of ACM Conference on Computer and Communications Security.
- [Venema 1996] Venema, Wietse. 1996. Murphy's law and computer security. <http://www.fish.com/security/murphy.html>
- [Watters 1996] Watters, Arron, Guido van Rossum, James C. Ahlstrom. 1996. Internet Programming with Python. NY, NY: Henry Hold and Company, Inc.
- [Wood 1985] Wood, Patrick H. and Stephen G. Kochan. 1985. *Unix System Security*. Indianapolis, Indiana: Hayden Books. ISBN 0-8104-6267-2.
- [Wreski 1998] Wreski, Dave. August 22, 1998. *Linux Security Administrator's Guide*. Version 0.98. <http://www.nic.com/~dave/SecurityAdminGuide/index.html>
- [Yoder 1998] Yoder, Joseph and Jeffrey Barcalow. 1998. Architectural Patterns for Enabling Application Security. PLoP '97 <http://st-www.cs.uiuc.edu/~hanmer/PLoP-97/Proceedings/yoder.pdf>
- [Zoebelein 1999] Zoebelein, Hans U. April 1999. The Internet Operating System Counter. <http://www.leb.net/hzo/ioscount>.

# Appendix A. History

Here are a few key events in the development of this document, starting from most recent events:

2000-05-24 David A. Wheeler

Switched to GNU's GFDL license, added more content.

2000-04-21 David A. Wheeler

Version 2.00 released, dated 21 April 2000, which switching the document's internal format from the Linuxdoc DTD to the DocBook DTD. Thanks to Jorge Godoy for helping me perform the transition.

2000-04-04 David A. Wheeler

Version 1.60 released; changed so that it now covers *both* Linux and Unix. Since most of the guidelines covered both, and many/most app developers want their apps to run on both, it made sense to cover both.

2000-02-09 David A. Wheeler

Noted that the document is now part of the Linux Documentation Project (LDP).

1999-11-29 David A. Wheeler

Initial version (1.0) completed and released to the public.

Note that a more detailed description of changes is available on-line in the "ChangeLog" file.

# Appendix B. Acknowledgements

*As iron sharpens iron, so one  
man sharpens another.  
Proverbs 27:17 (NIV)*

My thanks to the following people who kept me honest by sending me emails noting errors, suggesting areas to cover, asking questions, and so on. Where email addresses are included, they've been shrouded by prepending my "thanks." so bulk emailers won't easily get these addresses; inclusion of people in this list is *not* an authorization to send unsolicited bulk email to them.

- Neil Brown (thanks.neilb@cse.unsw.edu.au)
- Martin Douda (thanks.mad@students.zcu.cz)
- Jorge Godoy
- Scott Ingram (thanks.scott@silver.jhuapl.edu)
- Michael Kerrisk
- Doug Kilpatrick
- John Levon (moz@compsoc.man.ac.uk)
- Ryan McCabe (thanks.odin@numb.org)
- Paul Millar (thanks.paulm@astro.gla.ac.uk)
- Chuck Phillips (thanks.cdp@peakpeak.com)
- Martin Pool (thanks.mbp@humblebug.org.au)
- Eric S. Raymond (thanks.esr@snark.thyrsus.com)
- Marc Welz
- Eric Werme (thanks.werme@alpha.zk3.dec.com)

If you want to be on this list, please send me a constructive suggestion at [dwheeler@dwheeler.com](mailto:dwheeler@dwheeler.com) (<mailto:dwheeler@dwheeler.com>). If you send me a

constructive suggestion, but do *not* want credit, please let me know that when you send your suggestion, comment, or criticism; normally I expect that people want credit, and I want to give them that credit. My current process is to add contributor names to this list in the document, with more detailed explanation of their comment in the ChangeLog for this document (available on-line). Note that although these people have sent in ideas, the actual text is my own, so don't blame them for any errors that may remain. Instead, please send me another constructive suggestion.

# Appendix C. About the Documentation License

*A copy of the text of the edict was to be issued as law in every province and made known to the people of every nationality so they would be ready for that day.*

*Esther 3:14 (NIV)*

This document is Copyright (C) 1999-2000 David A. Wheeler. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License (FDL), Version 1.1 or any later version published by the Free Software Foundation; with the invariant sections being “About the Author”, with no Front-Cover Texts, and no Back-Cover texts. A copy of the license is included below in Appendix D.

These terms do permit mirroring by other web sites, but be *sure* to do the following:

- make sure your mirrors automatically get upgrades from the master site,
- clearly show the location of the master site (<http://www.dwheeler.com/secure-programs>), with a hypertext link to the master site, and
- give me (David A. Wheeler) credit as the author.

The first two points primarily protect me from repeatedly hearing about obsolete bugs. I do not want to hear about bugs I fixed a year ago, just because you are not properly mirroring the document. By linking to the master site, users can check and see if your mirror is up-to-date. I’m sensitive to the problems of sites which have very strong security requirements and therefore cannot risk normal connections to the Internet; if that describes your situation, at least try to meet the other points and try to occasionally sneakernet updates into your environment.

By this license, you may modify the document, but you can't claim that what you didn't write is yours (i.e., plagiarism) nor can you pretend that a modified version is identical to the original work. Modifying the work does not transfer copyright of the entire work to you; this is not a "public domain" work in terms of copyright law. See the license in Appendix D for details. If you have questions about what the license allows, please contact me. In most cases, it's better if you send your changes to the master integrator (currently David A. Wheeler), so that your changes will be integrated with everyone else's changes into the master copy.

# Appendix D. GNU Free Documentation License

Version 1.1, March 2000

Copyright © 2000

Free Software Foundation, Inc.  
59 Temple Place, Suite 330,  
Boston,  
MA  
02111-1307  
USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

## 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other written document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of

subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

## 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. The "Document" , below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you".

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (For example, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, whose contents can be viewed and edited directly and straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for



input to text formatters. A copy made in an otherwise Transparent file format whose markup has been designed to thwart or discourage subsequent modification by readers is not Transparent. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML designed for human modification. Opaque formats include PostScript, PDF, proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

## 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

## 3. COPYING IN QUANTITY

If you publish printed copies of the Document numbering more than 100, and the

Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a publicly-accessible computer-network location containing a complete Transparent copy of the Document, free of added material, which the general network-using public has access to download anonymously at no charge using public-standard network protocols. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

#### 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in

the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has less than five).
- C. State on the Title Page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section entitled "History", and its title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These

may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

- K. In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section as "Endorsements" or to conflict in title with any Invariant Section.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version .

## 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections entitled "History" in the various original documents, forming one section entitled "History"; likewise combine any sections entitled "Acknowledgements", and any sections entitled "Dedications". You must delete all sections entitled "Endorsements."

## 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding

verbatim copying of that document.

## 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, does not as a whole count as a Modified Version of the Document, provided no compilation copyright is claimed for the compilation. Such a compilation is called an "aggregate", and this License does not apply to the other self-contained works thus compiled with the Document, on account of their being thus compiled, if they are not themselves derivative works of the Document. If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one quarter of the entire aggregate, the Document's Cover Texts may be placed on covers that surround only the Document within the aggregate. Otherwise they must appear on covers around the whole aggregate.

## 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License provided that you also include the original English version of this License. In case of a disagreement between the translation and the original English version of this License, the original English version will prevail.

## 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

## 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation (<http://www.gnu.org/fsf/fsf.html>) may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/> (<http://www.gnu.org/copyleft/>).

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

### Addendum

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright © YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have no Invariant Sections, write "with no Invariant Sections" instead of saying which ones are invariant. If you have no Front-Cover Texts, write "no Front-Cover Texts" instead of "Front-Cover Texts being LIST"; likewise for Back-Cover Texts.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license,

such as the GNU General Public License (<http://www.gnu.org/copyleft/gpl.html>),  
to permit their use in free software.



# Appendix E. Endorsements

This version of the document is endorsed by the original author, David A. Wheeler, as a document that should improve the security of programs. when applied correctly. Modifications (including translations) must remove this appendix per the license agreement included above.

# Appendix F. About the Author

David A. Wheeler is an expert in computer security and has long specialized in development techniques for large and high-risk software systems. He has been involved in software development since the mid-1970s, and been involved with Unix and computer security since the early 1980s. His areas of knowledge include software safety, vulnerability analysis, inspections, Internet technologies, software-related standards (including POSIX), real-time software development techniques, and numerous computer languages (including Ada, C, C++, Perl, Python, and Java).

Mr. Wheeler is co-author and lead editor of the IEEE book *Software Inspection: An Industry Best Practice*, author of the book *Ada95: The Lovelace Tutorial*, and co-author of the *GNOME User's Guide*. He is also the author of many smaller papers and articles, including the *Linux Program Library HOWTO*.

Mr. Wheeler hopes that, by making this document available, other developers will make their software more secure. You can reach him by email at [dwheeler@dwheeler.com](mailto:dwheeler@dwheeler.com) (no spam please), and you can also see his web site at <http://www.dwheeler.com>.

