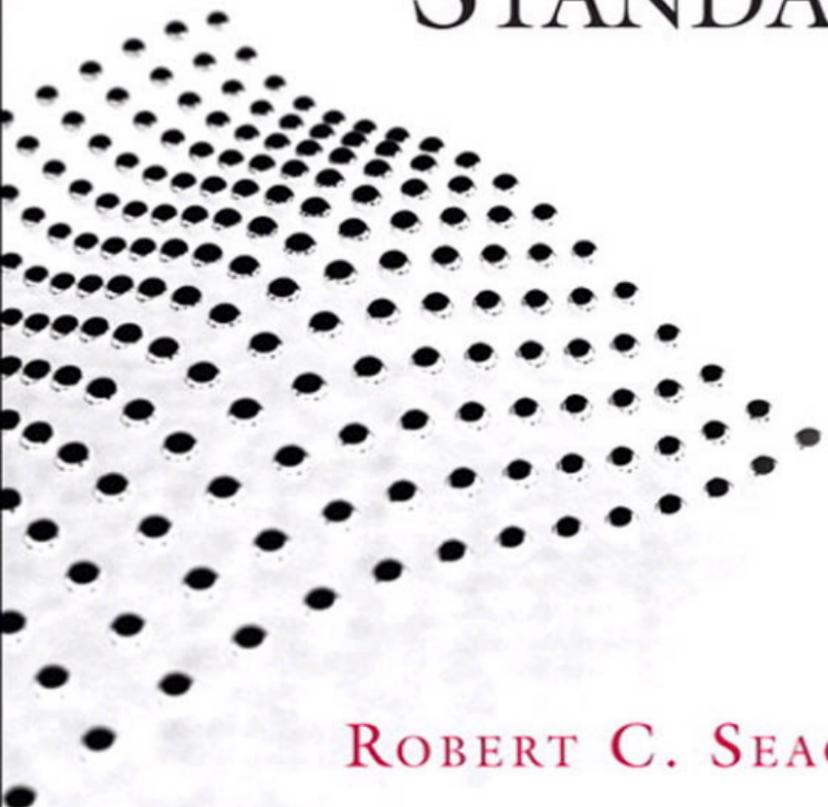


THE CERT® C SECURE CODING STANDARD



ROBERT C. SEACORD



**CarnegieMellon
Software Engineering Institute**

The SEI Series in Software Engineering

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

CMM, CMMI, Capability Maturity Model, Capability Maturity Modeling, Carnegie Mellon, CERT, and CERT Coordination Center are registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

ATAM; Architecture Tradeoff Analysis Method; CMM Integration; COTS Usage-Risk Evaluation; CURE; EPIC; Evolutionary Process for Integrating COTS Based Systems; Framework for Software Product Line Practice; IDEAL; Interim Profile; OAR; OCTAVE; Operationally Critical Threat, Asset, and Vulnerability Evaluation; Options Analysis for Reengineering; Personal Software Process; PLTP; Product Line Technical Probe; PSP; SCAMPI; SCAMPI Lead Appraiser; SCAMPI Lead Assessor; SCE; SEI; SEPG; Team Software Process; and TSP are service marks of Carnegie Mellon University.

Special permission to use material from the CERT Secure Coding Standards Website, © 2007 Carnegie Mellon University, in this publication is granted by the Software Engineering Institute.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside the U.S., please contact:
International Sales
international@pearsoned.com

Visit us on the Web: informit.com/aw

Library of Congress Cataloging-in-Publication Data

Seacord, Robert C.

The CERT C secure coding standard / Robert C. Seacord.
p. cm.

Includes bibliographical references and index.

ISBN 0-321-56321-2 (pbk. : alk. paper) 1. C (Computer program language) 2. Computer security. I. Title.

QA76.73.C15S4155 2008

005.8—dc22

2008030261

Copyright © 2009 Pearson Education, Inc.

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, write to:

Pearson Education, Inc.
Rights and Contracts Department
501 Boylston Street, Suite 900
Boston, MA 02116
Fax: (617) 671-3447

ISBN 13: 978-0-321-56321-7

ISBN 10: 0-321-56321-2

Text printed in the United States on recycled paper at Courier in Stoughton, Massachusetts.

First printing, October 2008.

Preface

An essential element of secure coding in the C programming language is a well-documented and enforceable coding standard. Coding standards encourage programmers to follow a uniform set of guidelines determined by the requirements of the project and organization rather than by the programmer's familiarity or preference. Once established, these standards can be used as a metric to evaluate source code (using manual or automated processes).

The CERT® C Secure Coding Standard provides guidelines for secure coding in the C programming language. The goal of these guidelines is to eliminate insecure coding practices and undefined behaviors that can lead to exploitable vulnerabilities. Developing code in compliance with this coding standard will result in higher quality systems that are robust and more resistant to attack.

This standard is supported by training available from the Software Engineering Institute (SEI) and other licensed partners and is a basis for the Global Information Assurance Certification (GIAC) Secure Software Programmer–C (GSSP-C) exam and certification.

■ The Demand for Secure Software

The Morris worm incident, which brought 10 percent of Internet systems to a halt in November 1988, resulted in a new and acute awareness of the need for secure software systems. Twenty years later, many security analysts, software developers, software users, and policymakers are asking the question, Why isn't software more secure?

The first problem is that the term *software security*, as it is used today, is meaningless. I have attempted to define this term, as have others, but there is no generally accepted definition. Why does this matter?

A variety of reasons are given for why software is not more secure: for example, the tools are inadequate, programmers lack sufficient training, and schedules are too short. But these are all solvable problems. The root cause of the issue lies elsewhere.

The reason software systems are not more secure is that there is no *demand* for secure software. In simple terms, if one vendor offers a product that has more features and better performance and is available today and another vendor offers a *secure* product that has fewer features and lesser performance and will be available in 6 months, there is really no question as to which product customers will buy, and vendors *know* this.

So why don't customers buy secure products? Again, this is because the word *secure* is meaningless in this context. Why would a customer pass up tangible benefits to buy a product that has an ill-defined and intangible property?

The problem is addressed by this coding standard. While developing code in compliance with this coding standard does not guarantee the security of a software system, it does tell you a great deal about the quality and security of the code. It tells you that the software was developed to a set of industry standard rules and recommendations that were developed by the leading experts in the field. It tells you that a tremendous amount of attention and effort went into producing code that is free from the common coding errors that have resulted in numerous vulnerabilities that have been reported to and published by the CERT Coordination Center (CERT/CC) over the past two decades. It tells you that the software developers who produced the code have done so with a real knowledge of the types of vulnerabilities that can exist and the exploits that can be used against them, and consequently have developed the software with a real security mindset in place.

So, the *small* problem we have set out to address in this book is to change the market dynamic for developing and purchasing software systems. By producing an *actionable and measurable* definition of software security for C language programs—compliance with the rules and recommendations in this standard—we have defined a mechanism by which customers can demand secure software systems and vendors can comply. Furthermore, the concept of a secure system now has *value* because the word *secure* has meaning.

■ Community Development Process

The CERT® C Secure Coding Standard was developed over a period of two and a half years as a community effort involving 226 contributors and reviewers.

The following development process was followed:

1. Rules and recommendations for a coding standard were solicited from the communities involved in the development and application of the C programming language, including the formal standard bodies responsible for the C language standard and user groups.

2. These rules and recommendations were edited by members of the CERT technical staff and industry experts for content and style on the CERT Secure Coding Standards wiki at www.securecoding.cert.org.
3. The user community reviewed and commented on the publicly posted content using threaded discussions and other communication tools. Drafts of this standard were reviewed at the London and Kona meetings by ISO/IEC WG14 and subjected to the scrutiny of the public, including members of the Association of C and C++ Users (ACCU) and the `comp.lang.c` newsgroup.

The Wiki versus This Book

Developing a secure coding standard on a wiki has many advantages. However, one disadvantage is that the content is constantly evolving. This is ideal if you want the latest information and are willing to entertain the possibility that a recent change has not yet been fully vetted. However, many software development organizations require a final document before they can commit to complying with a (fixed) set of rules and recommendations. This book serves that purpose as Version 1.0 of *The CERT® C Secure Coding Standard*.

Starting with the production of this book in June 2008, Version 1.0 and the wiki versions of the Secure Coding Standard began to diverge. Because both the C programming language and our knowledge of how to use it securely are still evolving, CERT will continue to evolve *The CERT® C Secure Coding Standard* on the Secure Coding wiki. These changes may then be incorporated into future, officially released versions of this standard.

Purpose

This book provides developers with *guidelines* for secure coding in the C programming language. These guidelines serve a variety of purposes. First, they enumerate common errors in C language programming that can lead to software defects, security flaws, and software vulnerabilities. These are all errors for which a conforming compiler is not required by the standard to issue a fatal diagnostic. In other words, the compiler will generate an executable, frequently without issuing any warnings, which can be shipped and deployed, and the resulting program may still contain flaws that make it vulnerable to attack.

Second, this coding standard provides recommendations for how to produce secure code. Failure to comply with these recommendations does not necessarily mean that the software is insecure, but if followed, these recommendations can be powerful tools in eliminating vulnerabilities from software.

Third, this coding standard identifies nonportable coding practices. Portability is not a strict requirement of security, but nonportable assumptions in code often result in vulnerabilities when code is ported to platforms for which these assumptions are no longer valid.

Rules

Guidelines are classified as either *rules* or *recommendations*. Guidelines are defined to be rules when all of the following conditions are met:

1. Violation of the coding practice is likely to result in a security flaw that may result in an exploitable vulnerability.
2. There is a denumerable set of conditions for which violating the coding practice is necessary to ensure correct behavior.
3. Conformance to the coding practice can be determined through automated analysis, formal methods, or manual inspection techniques.

Implementation of the secure coding rules defined in this standard are necessary (but not sufficient) to ensure the security of software systems developed in the C programming language. Figure P-1 shows how the 89 rules in this secure coding standard are categorized.

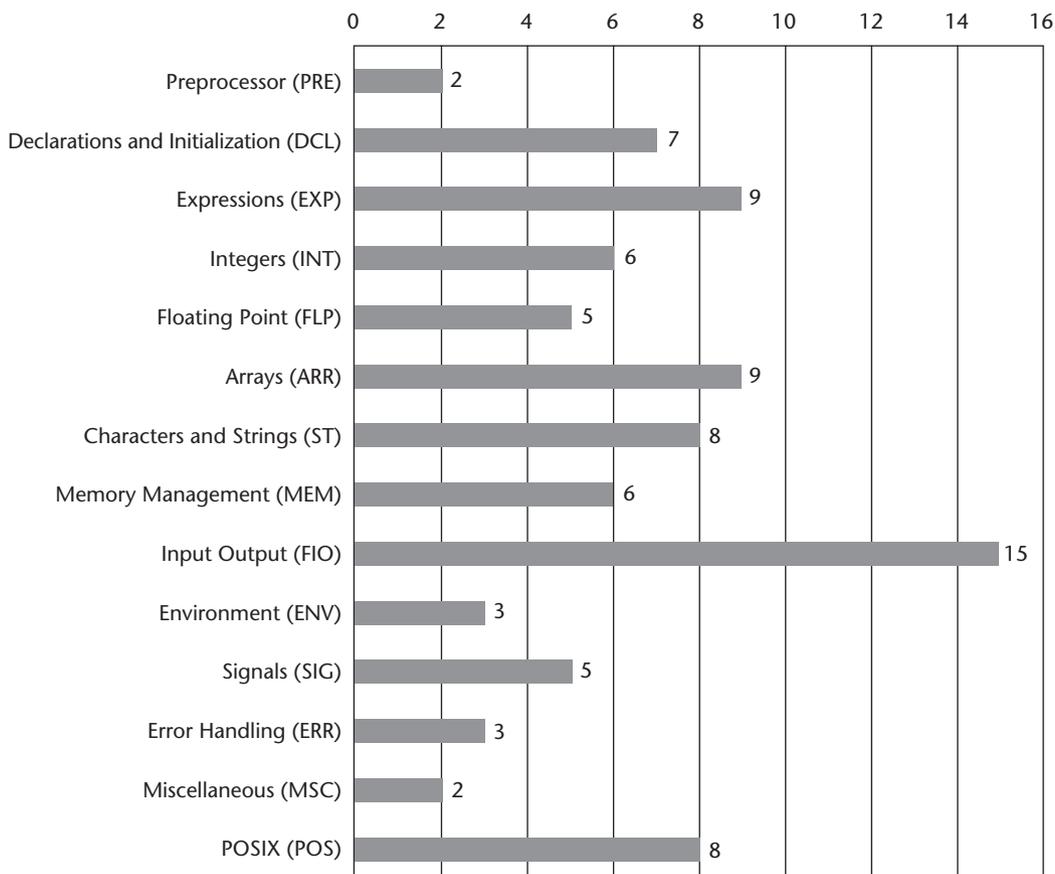


Figure P-1. CERT C Secure Coding rules

Recommendations

Guidelines are defined to be recommendations when all of the following conditions are met:

1. Application of the coding practice is likely to improve system security.
2. One or more of the requirements necessary for a coding practice to be considered a rule cannot be met.

The set of recommendations that a particular development effort adopts depends on the security requirements of the final software product. Projects with high-security requirements can dedicate more resources to security and consequently are likely to adopt a larger set of recommendations.

Figure P-2 shows how the 132 recommendations in this secure coding standard are categorized.

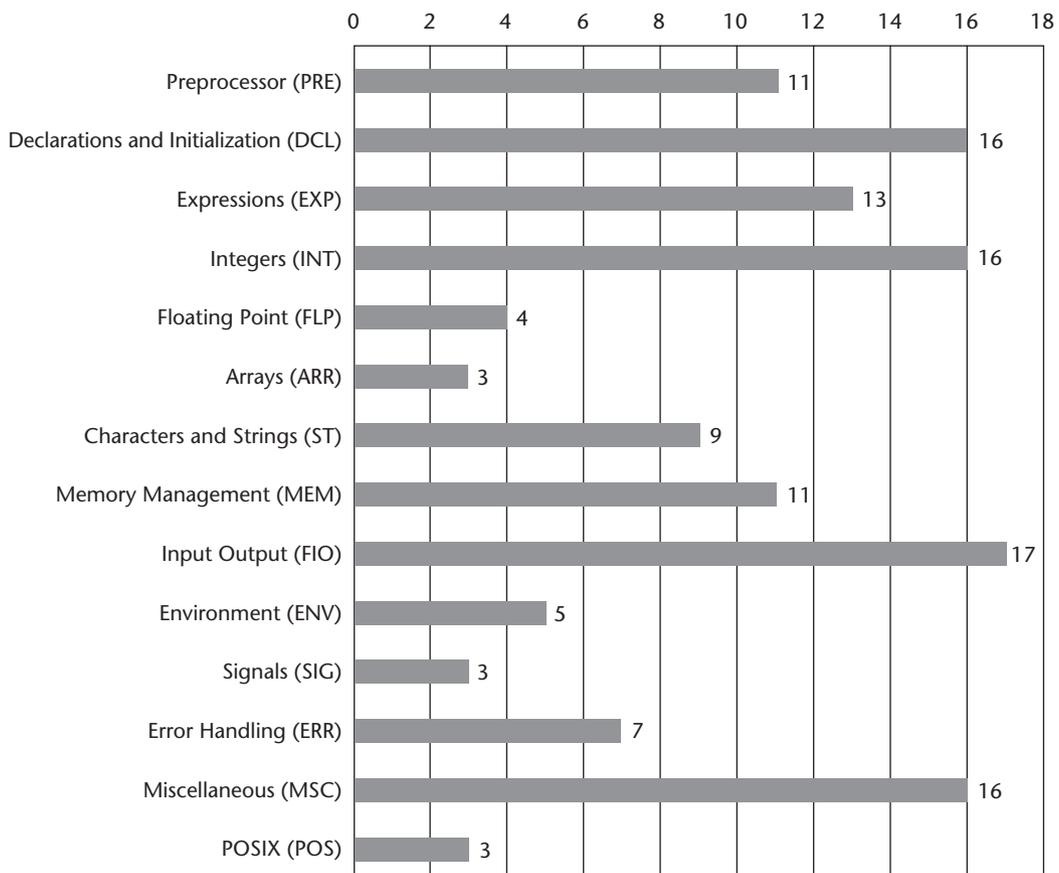


Figure P-2. CERT C Secure Coding recommendations

To ensure that the source code conforms to this secure coding standard, it is necessary to have measures in place that check for rules violations. The most effective means of achieving this is to use one or more static analysis tools. Where a rule cannot be checked by a tool, a manual review is required.

Both freely available and commercial source code analysis tools are available to automatically detect violations of CERT C Secure Coding Standard rules and recommendations, including Compass/ROSE, which has been developed by Lawrence Livermore National Laboratory and extended by CERT (www.rosecompiler.org).

■ Scope

The CERT® C Secure Coding Standard was developed specifically for versions of the C programming language defined in these publications:

- ISO/IEC 9899:1999, *Programming Languages—C*, Second Edition [ISO/IEC 9899:1999]
- Technical corrigenda TC1, TC2, and TC3
- ISO/IEC TR 24731-1, *Extensions to the C Library, Part I: Bounds-Checking Interfaces* [ISO/IEC TR 24731-1:2007]
- ISO/IEC PDTR 24731-2, *Extensions to the C Library, Part II: Dynamic Allocation Functions* [ISO/IEC PDTR 24731-2]

Most of the material included in this standard can also be applied to earlier versions of the C programming language.

Rules and recommendations included in this standard are designed to be operating system and platform independent. However, the best solutions to secure coding problems are often platform specific. In most cases, this standard provides appropriate compliant solutions for POSIX-compliant and Windows operating systems. In many cases, compliant solutions have also been provided for specific platforms such as Linux and OpenBSD. Occasionally, we also point out implementation-specific behaviors when these behaviors are of interest.

Rationale

A secure coding standard for the C programming language can create the highest value for the longest period of time by focusing on C99 and the relevant post-C99 technical reports.

In addition, because considerably more money and effort is devoted to developing new code than maintaining existing code, the highest return on investment comes from influencing programmers who are developing new code [Seacord 03]. Maintaining existing code is still an important concern, however.

The C standard (C99) documents existing practice where possible [ISO/IEC 9899:1999]. That is, most features must be tested in an implementation before being included in the standard. *The CERT® C Secure Coding Standard* has a different purpose. When existing practice serves this purpose, that is fine, but the goal is to create a new set of best practices, and that includes introducing some concepts that are not yet widely known. To put it a different way, the CERT C secure coding guidelines are attempting to drive change rather than just document it.

For example, the C library technical report, part 1 (TR 24731-1), is gaining support, but at present it is implemented by only a few vendors. It introduces functions such as `memcpy_s()`, which serve the purpose of security by adding the destination buffer size to the API. A forward-looking document could not reasonably ignore such functions simply because they are not yet widely implemented.

C99 is more widely implemented than TR 24731-1, but even if it were not yet, it is the direction in which the industry is moving. Developers of new C code, especially, need guidance that is usable on and makes the best use of the compilers and tools that are now being developed and will be supported into the future.

Some vendors have extensions to C, and some have implemented only part of the C standard before stopping development. Consequently, it is not possible to back up and only discuss C95 or C90. The vendor support equation is too complicated to draw a line and say that a certain compiler supports exactly a certain standard. Whatever demarcation point is selected, different vendors are on opposite sides of it for different parts of the language. Supporting all possibilities would require testing the cross product of each compiler with each language feature. Consequently, a recent demarcation point was selected so that the rules and recommendations defined by the standard will be applicable for as long as possible. As a result of the variations in support, source code portability is enhanced when the programmer uses only the features specified by C90. This is one of many tradeoffs between security and portability inherent to C language programming.

The value of forward-looking information increases with time before it starts to decrease. The value of backward-looking information starts to decrease immediately.

For all of these reasons, the priority of this standard is to support new code development using C99 and the post-C99 technical reports. A close-second priority is supporting remediation of old code using C99 and the technical reports.

This standard does try to make contributions to support older compilers when these contributions can be significant and doing so does not compromise other priorities. The intent is not to capture all deviations from the standard but only a few important ones.

Issues Not Addressed

There are a number of issues not addressed by this secure coding standard.

- **Coding Style.** Coding style issues are subjective, and it has proven impossible to develop a consensus on appropriate style guidelines. Consequently, this standard does not require any particular coding style to be enforced but only that the user define style guidelines and apply those guidelines consistently. The easiest way to consistently apply a coding style is with the use of a code formatting tool. Many interactive development environments (IDEs) provide such capabilities.
- **Tools.** As a federally funded research and development center (FFRDC), the SEI is not in a position to recommend particular vendors or tools to enforce these guidelines. The user of this document is free to choose tools, and vendors are encouraged to provide tools to enforce this standard.
- **Controversial Rules.** In general, the CERT secure coding standards try to avoid the inclusion of controversial rules that lack a broad consensus.

■ Who Should Read This Book

The CERT® C Secure Coding Standard is primarily intended for developers of C language programs. While security is important for Internet-facing systems, for example, it is also important for any software component that may be included or deployed as part of a secure software system. With systems increasingly being composed of software components, or even other systems, it is difficult to identify situations in which software is guaranteed not to be used in another context, which perhaps has more stringent security requirements.

This book is also useful for C language programmers who don't realize they are interested in security, as most of these guidelines have practical applications for achieving other quality attributes such as safety, reliability, dependability, robustness, availability, and maintainability.

While not intended for C++ programmers, this book may be of some value because the vast majority of issues identified for C language programs are also issues in C++ programs, although in many cases the solutions are different.

■ How This Book Is Organized

This book is organized into an introductory chapter, thirteen chapters containing guidelines in specific topic areas, and an appendix containing POSIX guidelines to demonstrate

how this secure coding standard can be customized for particular environments. The POSIX appendix is nonnormative and not a prescriptive part of the standard.

Most guidelines have a consistent structure. Each guideline in this standard has a unique *identifier*, which is included in the title. The title of the guidelines and the introductory paragraphs define the rule or recommendation. This is typically followed by one or more pairs of *noncompliant code examples* and *compliant solutions*. Each guideline also includes a *risk assessment* and a list of appropriate *references* (where applicable). Guidelines may also include a table of *related vulnerabilities*.

Guideline Identifiers

Guideline identifiers consist of three parts:

- a three-letter mnemonic representing the section of the standard
- a two-digit numeric value in the range of 00 to 99
- the letter *C* indicating that this is a C language guideline

The three-letter mnemonic can be used to group similar guidelines and to indicate to which category a guideline belongs.

The numeric value is used to give each guideline a unique identifier. Numeric values in the range of 00 to 29 are reserved for recommendations, while values in the range of 30 to 99 are reserved for rules.

Noncompliant Code Examples and Compliant Solutions

Noncompliant code examples are examples of insecure code that violate the guideline under discussion. It is important to note that these are only examples, and eliminating all occurrences of the example does not necessarily mean that your code is now compliant with the guideline.

The noncompliant code examples are typically followed by compliant solutions, which show how the noncompliant code example can be reimplemented in a secure, compliant manner. Except where noted, noncompliant code examples should only contain violations of the rule under discussion. Compliant solutions should comply with all secure coding rules but may occasionally fail to comply with a recommendation.

Risk Assessment

Each guideline contains a risk assessment section, which attempts to quantify and qualify the risk of violating each guideline. This information is intended primarily for remediation projects to help prioritize repairs, as it is assumed that new code will be developed in conformance with the entire standard.

Each rule and recommendation has an assigned priority. Priorities are assigned using a metric based on Failure Mode, Effects, and Criticality Analysis (FMECA) [IEC 60812]. Three values are assigned for each rule on a scale of 1 to 3 for

- Severity: How serious are the consequences of the rule being ignored?
 - 1 = low (denial-of-service attack, abnormal termination)
 - 2 = medium (data integrity violation, unintentional information disclosure)
 - 3 = high (run arbitrary code)
- Likelihood: How likely is it that a flaw introduced by ignoring the rule could lead to an exploitable vulnerability?
 - 1 = unlikely
 - 2 = probable
 - 3 = likely
- Remediation cost: How expensive is it to comply with the rule?
 - 1 = high (manual detection and correction)
 - 2 = medium (automatic detection and manual correction)
 - 3 = low (automatic detection and correction)

The three values are multiplied together for each rule. This product provides a measure that can be used in prioritizing the application of the rules. These products range from 1 to 27. Rules and recommendations with a priority in the range of 1 to 4 are level 3 rules, 6 to 9 are level 2, and 12 to 27 are level 1. As a result, it is possible to claim level 1, level 2, or complete compliance (level 3) with a standard by implementing all rules in a level, as shown in Figure P-3.

Recommendations are not compulsory, and risk assessments are provided for information purposes only.

References

Guidelines include frequent references to the vulnerability notes in the CERT/CC Vulnerability Notes Database [CERT/CC VND], CWE IDs in MITRE's Common Weakness Enumeration (CWE) [MITRE 07], and CVE numbers from MITRE's Common Vulnerabilities and Exposures (CVE) [CVE].

You can create a unique URL to get more information on any of these topics by appending the relevant ID to the end of a fixed string. For example, to find more information about

- VU#551436, "Mozilla Firefox SVG viewer vulnerable to integer overflow," you can append 551436 to <https://www.kb.cert.org/vulnotes/id/> and enter the resulting URL in your browser: <https://www.kb.cert.org/vulnotes/id/551436>

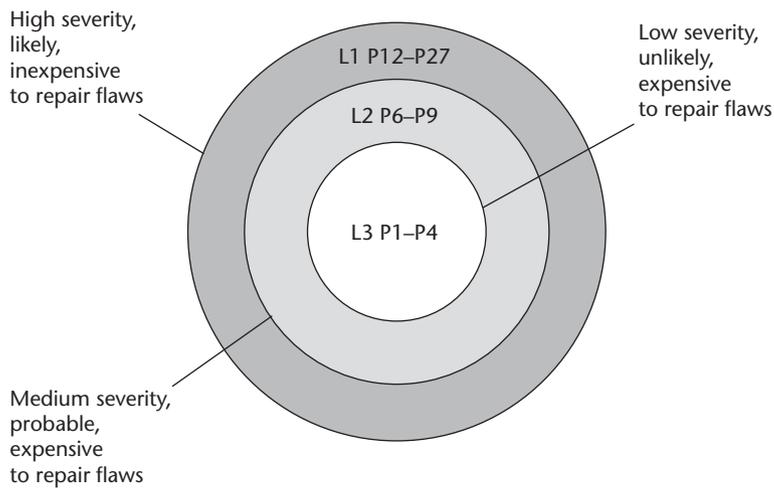


Figure P-3. Priorities and levels

- CWE ID 192, “Integer Coercion Error,” you can append 192.html to <http://cwe.mitre.org/data/definitions/> and enter the resulting URL in your browser: <http://cwe.mitre.org/data/definitions/192.html>
- CVE-2006-1174, you can append CVE-2006-1174 to <http://cve.mitre.org/cgi-bin/cve-name.cgi?name=> and enter the resulting URL in your browser: [http://cve.mitre.org/cgi-bin/cvename.cgi?name= CVE-2006-1174](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-1174)

Guidelines are frequently correlated with language vulnerabilities in *Information Technology—Programming Languages—Guidance to Avoiding Vulnerabilities in Programming Languages through Language Selection and Use* [ISO/IEC PDTR 24772].

Related Vulnerabilities

Rules and recommendations linked to violations of actual vulnerabilities published in the CERT/CC Vulnerability Notes Database are shown in sections marked “Related Vulnerabilities” and are presented in table format, as in this example:

Metric	ID	Date Public	Name
1.62	VU#606700	03/19/2007	file integer overflow vulnerability
2.06	VU#559444	03/13/2007	Apple Mac OS X ImageIO integer overflow vulnerability

Metric	ID	Date Public	Name
22.22	VU#551436	02/23/2007	Mozilla Firefox SVG viewer vulnerable to integer overflow
0	VU#162289	04/17/2006	C compilers may silently discard some wraparound checks

New links are continually added. To find the latest list of related vulnerabilities, enter the following URL:

<https://www.kb.cert.org/vulnotes/bymetric?searchview&query=FIELD+KEYWORDS+contains+XXXNN-X>

where XXXNN-X is the ID of the rule or recommendation for which you are searching.

These tables consist of four fields: *metric*, *ID*, *date public*, and *name*.

Vulnerability Metric. The CERT vulnerability metric value is a number between 0 and 180 that assigns an approximate severity to the vulnerability. This value incorporates several elements:

- Is information about the vulnerability widely available or known?
- Is the vulnerability being exploited in incidents reported to CERT or other incident response teams?
- Is the Internet infrastructure (e.g., routers, name servers, critical Internet protocols) at risk because of this vulnerability?
- How many systems on the Internet are at risk from this vulnerability?
- What is the impact of exploiting the vulnerability?
- How easy is it to exploit the vulnerability?
- What are the preconditions required to exploit the vulnerability?

Because the questions are answered with approximate values based on the judgment of vulnerability analysts and may differ significantly from one site to another, you should not rely too heavily on the metric for prioritizing response to vulnerabilities. Rather, this metric may be useful for separating the serious vulnerabilities from the larger number of less severe vulnerabilities described in the database. Because the questions are not all weighted equally, the resulting score is not linear (that is, a vulnerability with a metric of 40 is not twice as severe as one with a metric of 20).

An alternative vulnerability severity metric is the Common Vulnerability Scoring System (CVSS) [Mell 07].

Vulnerability ID. Vulnerability ID numbers are assigned at random to uniquely identify a vulnerability. These IDs are four to six digits long and are usually prefixed with VU# to mark them as vulnerability IDs.

Date Public. This is the date on which the vulnerability was first publicly disclosed. Usually this date is when the vulnerability note was first published, when an exploit was first discovered, when the vendor first distributed a patch publicly, or when a description of the vulnerability was posted to a public mailing list. By default, this date is set to the vulnerability note publication date.

Vulnerability Name. The vulnerability name is a short description that summarizes the nature of the problem and the affected software product. While the name may include a clause describing the impact of the vulnerability, most names are focused on the nature of the defect that caused the problem to occur.

Type-generic macros may also be used, for example, to swap two variables of any type, provided they are of the same type.

PRE00-EX5: Macro parameters exhibit call-by-name semantics, whereas functions are call-by-value. Macros must be used in cases where call-by-name semantics are required.

Risk Assessment

Improper use of macros may result in undefined behavior.

Recommendation	Severity	Likelihood	Remediation Cost	Priority	Level
PRE00-C	medium	unlikely	medium	P4	L3

References

- [FSF 05] Section 5.34, “An Inline Function Is as Fast as a Macro”
- [Dewhurst 02] Gotcha #26, “#define Pseudofunctions”
- [ISO/IEC 9899:1999] Section 6.7.4, “Function Specifiers”
- [ISO/IEC PDTR 24772] “NMP Pre-processor Directives”
- [Kettlewell 03]
- [MISRA 04] Rule 19.7
- [Summit 05] Question 10.4

■ PRE01-C. Use parentheses within macros around parameter names

Parenthesize all parameter names found in macro definitions. See also PRE00-C, “Prefer inline or static functions to function-like macros,” and PRE02-C, “Macro replacement lists should be parenthesized.”

Noncompliant Code Example

This CUBE() macro definition is noncompliant because it fails to parenthesize the parameter names.

```
#define CUBE(I) (I * I * I)
```

As a result, the invocation

```
int a = 81 / CUBE(2 + 1);
```

expands to

```
int a = 81 / (2 + 1 * 2 + 1 * 2 + 1); /* evaluates to 11 */
```

which is clearly not the desired result.

Compliant Solution

Parenthesizing all parameter names in the CUBE() macro allows it to expand correctly (when invoked in this manner).

```
#define CUBE(I) ( (I) * (I) * (I) )
int a = 81 / CUBE(2 + 1);
```

Exceptions

PRE01-EX1: When the parameter names are surrounded by commas in the replacement text, regardless of how complicated the actual arguments are, there is no need for parenthesizing the macro parameters. Because commas have lower precedence than any other operator, there is no chance of the actual arguments being parsed in a surprising way. Comma separators, which separate arguments in a function call, also have lower precedence than other operators, although they are technically different from comma operators.

```
#define F00(a, b, c) bar(a, b, c)
/* ... */
F00(arg1, arg2, arg3);
```

PRE01-EX2: Macro parameters cannot be individually parenthesized when concatenating tokens using the ## operator, converting macro parameters to strings using the # operator, or concatenating adjacent string literals. The JOIN() macro below concatenates both arguments to form a new token. The SHOW() macro converts the single argument into a string literal, which is then concatenated with the adjacent string literal to form the format specification in the call to printf().

```
#define JOIN(a, b) (a ## b)
#define SHOW(a) printf("#a " = %d\n", a)
```

See PRE05-C, “Understand macro replacement when concatenating tokens or performing stringification,” for more information on using the ## operator to concatenate tokens.

Risk Assessment

Failing to parenthesize the parameter names in a macro can result in unintended program behavior.

Recommendation	Severity	Likelihood	Remediation Cost	Priority	Level
PRE01-C	medium	probable	low	P12	L1

References

- [ISO/IEC 9899:1999] Section 6.10, “Preprocessing Directives,” and Section 5.1.1, “Translation Environment”
- [ISO/IEC PDTR 24772] “JCW Operator Precedence/Order of Evaluation”
- [MISRA 04] Rule 19.1
- [Plum 85]
- [Summit 05] Question 10.1

■ PRE02-C. Macro replacement lists should be parenthesized

Macro replacement lists should be parenthesized to protect any lower-precedence operators from the surrounding expression. See also PRE00-C, “Prefer inline or static functions to function-like macros,” and PRE01-C, “Use parentheses within macros around parameter names.”

Noncompliant Code Example

This CUBE() macro definition is noncompliant because it fails to parenthesize the replacement list.

```
#define CUBE(X) (X) * (X) * (X)
int i = 3;
int a = 81 / CUBE(i);
```

As a result, the invocation

```
int a = 81 / CUBE(i);
```

expands to

```
int a = 81 / i * i * i;
```

which evaluates as

```
int a = ((81 / i) * i) * i; /* evaluates to 243 */
```

which is not the desired behavior.

existing destination file in FIO10-C, “Take care when using the `rename()` function,” for an example of this exception.

EXP12-EX2: If a function cannot fail or if the return value cannot signify an error condition, the return value may be ignored. Such functions should be added to a white list when automatic checkers are used.

```
strcpy(dst, src);
```

Risk Assessment

Failure to handle error codes or other values returned by functions can lead to incorrect program flow and violations of data integrity.

Recommendation	Severity	Likelihood	Remediation Cost	Priority	Level
EXP12-C	medium	unlikely	medium	P4	L3

References

- [ISO/IEC 9899:1999] Section 6.8.3, “Expression and Null Statements”
- [ISO/IEC PDTR 24772] “CSJ Passing Parameters and Return Values”

■ EXP30-C. Do not depend on order of evaluation between sequence points

Evaluation of an expression may produce side effects. At specific points during execution called sequence points, all side effects of previous evaluations have completed and no side effects of subsequent evaluations have yet taken place.

According to C99, Section 6.5:

Between the previous and next sequence point an object can only have its stored value modified once by the evaluation of an expression. Additionally, the prior value can be read only to determine the value to be stored.

This requirement must be met for each allowable ordering of the subexpressions of a full expression; otherwise the behavior is undefined.

This rule means that statements such as

```
i = i + 1;
a[i] = i;
```

are allowed, while statements like

```
/* i is modified twice between sequence points */  
i = ++i + 1;  
  
/* i is read other than to determine the value to be stored */  
a[i++] = i;
```

are not.

Noncompliant Code Example

Programs cannot safely rely on the order of evaluation of operands between sequence points. In this noncompliant code example, the order of evaluation of the operands to the + operator is unspecified.

```
a = i + b[++i];
```

If *i* was equal to 0 before the statement, the statement may result in the following outcome:

```
a = 0 + b[1];
```

Or it may result in the following outcome:

```
a = 1 + b[1];
```

Compliant Solution

These examples are independent of the order of evaluation of the operands and can be interpreted in only one way.

```
++i;  
a = i + b[i];
```

Or alternatively:

```
a = i + b[i+1];  
++i;
```

Noncompliant Code Example

The order of evaluation for function arguments is unspecified.

```
func(i++, i);
```

The call to `func()` has undefined behavior because there are no sequence points between the argument expressions. The first (left) argument expression reads the value of `i` (to determine the value to be stored) and then modifies `i`. The second (right) argument expression reads the value of `i` between the same pair of sequence points as the first argument, but not to determine the value to be stored in `i`. This additional attempt to read the value of `i` has undefined behavior.

Compliant Solution

This solution is appropriate when the programmer intends for both arguments to `func()` to be equivalent.

```
i++;  
func(i, i);
```

This solution is appropriate when the programmer intends for the second argument to be one greater than the first.

```
j = i++;  
func(j, i);
```

Risk Assessment

Attempting to modify an object multiple times between sequence points may cause that object to take on an unexpected value. This can lead to unexpected program behavior.

Rule	Severity	Likelihood	Remediation Cost	Priority	Level
EXP30-C	medium	probable	medium	P8	L2

References

- [ISO/IEC 9899:1999] Section 5.1.2.3, “Program Execution,” Section 6.5, “Expressions,” and Annex C, “Sequence Points”
- [ISO/IEC PDTR 24772] “JCW Operator Precedence/Order of Evaluation” and “Side-Effects and Order of Evaluation [SAM]”
- [MISRA 04] Rule 12.1
- [Summit 05] Questions 3.1, 3.2, 3.3, 3.3b, 3.7, 3.8, 3.9, 3.10a, 3.10b, and 3.11
- [Saks 07a]

Exceptions

INT34-EX1: Unsigned integers can exhibit modulo behavior as long as the variable declaration is clearly commented as supporting modulo behavior, and each operation on that integer is also clearly commented as supporting modulo behavior.

If the integer exhibiting modulo behavior contributes to the value of an integer not marked as exhibiting modulo behavior, the resulting integer must obey this rule.

Risk Assessment

Improper range checking can lead to buffer overflows and the execution of arbitrary code by an attacker.

Rule	Severity	Likelihood	Remediation Cost	Priority	Level
INT34-C	high	probable	medium	P12	L1

References

- [Dowd 06] Chapter 6, “C Language Issues”
- [ISO/IEC 03] Section 6.5.7, “Bitwise Shift Operators”
- [ISO/IEC 9899:1999] Section 6.5.7, “Bitwise Shift Operators”
- [ISO/IEC PDTR 24772] “XYX Wrap-around Error”
- [Seacord 05a] Chapter 5, “Integers”
- [Viega 05] Section 5.2.7, “Integer Overflow”
- A test program for this rule is available at www.securecoding.cert.org/confluence/download/attachments/4385/leftshift.cpp.

■ INT35-C. Evaluate integer expressions in a larger size before comparing or assigning to that size

If an integer expression is compared to, or assigned to, a larger integer size, that integer expression should be evaluated in that larger size by explicitly casting one of the operands.

Noncompliant Code Example

This noncompliant code example is noncompliant on systems where `size_t` is an unsigned 32-bit value and `long long` is a 64-bit value. In this example, the programmer tests for wrapping by comparing `SIZE_MAX` to `length + BLOCK_HEADER_SIZE`. Because

`length` is declared as `size_t`, however, the addition is performed as a 32-bit operation and can result in wrapping. The comparison with `SIZE_MAX` will always test false. If wrapping occurs, `malloc()` will allocate insufficient space for `mBlock`, which can lead to a subsequent buffer overflow.

```
enum { BLOCK_HEADER_SIZE = 16 };

void *AllocateBlock(size_t length) {
    struct memBlock *mBlock;

    if (length + BLOCK_HEADER_SIZE > (unsigned long long)SIZE_MAX)
        return NULL;
    mBlock = (struct memBlock *)malloc(
        length + BLOCK_HEADER_SIZE
    );
    if (!mBlock) return NULL;

    /* fill in block header and return data portion */

    return mBlock;
}
```

Some compilers will diagnose this condition.

Compliant Solution (Upcast)

In this compliant solution, the `length` operand is upcast to `unsigned long long`, ensuring that the addition takes place in this size.

```
enum { BLOCK_HEADER_SIZE = 16 };

void *AllocateBlock(size_t length) {
    struct memBlock *mBlock;

    if ((unsigned long long)length + BLOCK_HEADER_SIZE > SIZE_MAX) {
        return NULL;
    }
    mBlock = (struct memBlock *)malloc(
        length + BLOCK_HEADER_SIZE
    );
    if (!mBlock) return NULL;

    /* fill in block header and return data portion */

    return mBlock;
}
```

This test for wrapping is effective only when the `sizeof(unsigned long long) > sizeof(size_t)`. If both `size_t` and `unsigned long long` types are represented as a 64-bit unsigned value, the result of the addition operation may not be representable as an `unsigned long long` value.

Compliant Solution (Rearrange Expression)

In this compliant solution, `length` is subtracted from `SIZE_MAX`, ensuring that wrapping cannot occur (see INT30-C, “Ensure that unsigned integer operations do not wrap”).

```
enum { BLOCK_HEADER_SIZE = 16 };

void *AllocateBlock(size_t length) {
    struct memBlock *mBlock;

    if (SIZE_MAX - length < BLOCK_HEADER_SIZE) return NULL;
    mBlock = (struct memBlock *)malloc(
        length + BLOCK_HEADER_SIZE
    );
    if (!mBlock) return NULL;

    /* fill in block header and return data portion */

    return mBlock;
}
```

Noncompliant Code Example

In this noncompliant code example, the programmer attempts to prevent wrapping by allocating an `unsigned long long` integer called `alloc` and assigning it the result from `cBlocks * 16`.

```
void* AllocBlocks(size_t cBlocks) {
    if (cBlocks == 0) return NULL;
    unsigned long long alloc = cBlocks * 16;
    return (alloc < UINT_MAX) ? malloc(cBlocks * 16) : NULL;
}
```

There are two separate problems with this noncompliant code example. The first problem is that this code assumes an implementation where `unsigned long long` has a least four more bits than `size_t`. The second problem, assuming an implementation where `size_t` is a 32-bit value and `unsigned long long` is represented by a 64-bit value, is that to be compliant with C99, multiplying two 32-bit numbers in this context must yield a 32-bit result. Any wrapping resulting from this multiplication will remain undetected by this code, and the expression `alloc < UINT_MAX` will always be true.

Compliant Solution

In this compliant solution, the `cBlocks` operand is upcast to `unsigned long long`, ensuring that the multiplication takes place in this size.

```
static_assert(
    CHAR_BIT * sizeof(unsigned long long) >=
    CHAR_BIT * sizeof(size_t) + 4,
    "Unable to detect wrapping after multiplication"
);

void* AllocBlocks(size_t cBlocks) {
    if (cBlocks == 0) return NULL;
    unsigned long long alloc = (unsigned long long)cBlocks * 16;
    return (alloc < UINT_MAX) ? malloc(cBlocks * 16) : NULL;
}
```

Note that this code does not prevent wrapping unless the `unsigned long long` type is at least four bits larger than `size_t`.

Risk Assessment

Failure to cast integers before comparing or assigning them to a larger integer size can result in software vulnerabilities that can allow the execution of arbitrary code by an attacker with the permissions of the vulnerable process.

Rule	Severity	Likelihood	Remediation Cost	Priority	Level
INT35-C	high	likely	medium	P18	L1

References

- [Dowd 06] Chapter 6, “C Language Issues”
- [ISO/IEC 9899:1999] Section 6.3.1, “Arithmetic Operands”
- [ISO/IEC PDTR 24772] “FLC Numeric Conversion Errors”
- [MITRE 07] CWE ID 681, “Incorrect Conversion between Numeric Types,” and CWE ID 190, “Integer Overflow (Wrap or Wraparound)”
- [Seacord 05a] Chapter 5, “Integer Security”

References

- [Banahan 03] Section 5.3, “Pointers,” and Section 5.7, “Expressions Involving Pointers”
- [ISO/IEC 9899:1999] Section 6.5.6, “Additive Operators”
- [MITRE 07] CWE ID 469, “Use of Pointer Subtraction to Determine Size”
- [VU#162289]

■ ARR38-C. Do not add or subtract an integer to a pointer if the resulting value does not refer to a valid array element

Do not add or subtract an integer to a pointer if the resulting value does not refer to an element within the array (or to the nonexistent element just after the last element of the array). According to C99, Section 6.5.6:

If both the pointer operand and the result point to elements of the same array object, or one past the last element of the array object, the evaluation shall not produce an overflow; otherwise, the behavior is undefined.

If the pointer resulting from the addition (or subtraction) is outside of the bounds of the array, an overflow has occurred and the result is undefined.

Noncompliant Code Example

In this noncompliant code example, a pointer is set to reference the start of an array. Array elements are accessed sequentially within the for loop. The array pointer `ip` is incremented on each iteration.

```
int ar[20];
int *ip;

for (ip = &ar[0]; ip < &ar[21]; ip++) {
    *ip = 0;
}
```

C99 guarantees that it is permissible to use the address of `ar[20]` even though no such element exists. However, in this noncompliant code example, the bound of the array is incorrectly specified, and consequently, the reference to `&ar[21]` constitutes undefined behavior. On the final iteration of the loop, the expression `ip++` (which adds 1 to `ip`) will also overflow.

This code also suffers from using “magic numbers,” described in DCL06-C, “Use meaningful symbolic constants to represent literal values in program logic.” When replacing the numbers with constants, a developer is likely to catch the invalid array bounds in the `for` statement.

Compliant Solution

This compliant solution fixes the problem from the previous noncompliant code example by using the common idiom `sizeof(ar)/sizeof(ar[0])` to determine the actual number of elements in the array. This idiom works only when the definition of the array is visible (see ARR01-C, “Do not apply the `sizeof` operator to a pointer when taking the size of an array”).

```
int ar[20];
int *ip;

for (ip = &ar[0]; ip < &ar[sizeof(ar)/sizeof(ar[0])]; ip++) {
    *ip = 0;
}
```

C99 guarantees that it is permissible to use the address of `ar[sizeof(ar)/sizeof(ar[0])]` even though no such element exists. This allows you to use this address for checks in loops like the one above. The guarantee extends to one element beyond the end of an array and no further [Banahan 03].

Noncompliant Code Example

Pointer arithmetic can result in undefined behavior if the pointer operand and the resulting pointer do not refer to the same array object (or one past the last element of the array object). Compiler implementations are provided broad latitude by the standard in how to deal with undefined behavior (see MSC15-C, “Do not depend on undefined behavior”), including ignoring the situation completely with unpredictable results.

In this noncompliant code example, the programmer is trying to determine if a pointer added to a length will wrap around the end of memory.

```
char *buf;
size_t len = 1 << 30;

/* Check for overflow */
if (buf + len < buf) {
    len = -(uintptr_t)buf-1;
}
```

This code resembles the test for wraparound from the `sprint()` function as implemented for the Plan 9 operating system. If `buf + len < buf` evaluates to true, `len` is assigned the remaining space minus 1 byte. However, because the expression `buf + len < buf` constitutes undefined behavior, compilers can assume this condition will never occur and optimize out the entire conditional statement. In gcc version 4.2 and later, for example, code that performs checks for wrapping that depend on undefined behavior (such as the code in this noncompliant code example) are optimized away; no object code to perform the check appears in the resulting executable program [VU#162289]. This is of special concern because it often results in the silent elimination of code that was inserted to provide a safety or security check. For gcc version 4.2.4 and later, this optimization may be disabled with the `-fno-strict-overflow` option.

Compliant Solution (Linear Address Space)

In this compliant solution, references to `buf` are cast to `uintptr_t`. The `uintptr_t` type is an unsigned integer type with the property that any valid pointer to `void` can be converted to this type, then converted back to pointer to `void`, and the result will compare equal to the original pointer. Because it is an unsigned type, C99 guarantees that it has modulo behavior. Alternatively, developers can use `size_t` on platforms that do not provide the `uintptr_t` type.

```
char *buf;
size_t len = 1 << 30;

/* Check for overflow */
if ((uintptr_t)buf+len < (uintptr_t)buf) {
    len = -(uintptr_t)buf-1;
}
```

This compliant solution works on architectures that provide a linear address space and the `uintptr_t` value is a byte address. The latter is not guaranteed, although it is usually the case. Some word-oriented machines are likely to produce a word address with the high-order bits used as a byte selector, in which case this solution will fail. Consequently, this is not a portable solution.

This same compliant solution can be implemented without wrapping:

```
char *buf;
size_t len = 1 << 30;

/* Check for overflow */
if (UINTPTR_MAX - len < (uintptr_t)buf) {
    len = -(uintptr_t)buf-1;
}
```

Noncompliant Code Example

Another interesting case is shown in this noncompliant code example. The expression `buf + n` may wrap for large values of `n`, resulting in undefined behavior.

```
int process_array(char *buf, size_t n) {
    return buf + n < buf + 100;
}
```

This is an example of how optimization may actually help improve security. When compiled using GCC 4.3.0 with the `-O2` option, for example, the expression `buf + n < buf + 100` is optimized to `n < 100`, eliminating the possibility of wrapping. This code example is still noncompliant because it is not safe to rely on compiler optimizations for security.

Compliant Solution

In this compliant solution, the “optimization” is performed by hand.

```
int process_array(char *buf, size_t n) {
    return n < 100;
}
```

Risk Assessment

If adding or subtracting an integer to a pointer results in a reference to an element outside the array or one past the last element of the array object, the behavior is undefined but frequently leads to a buffer overflow, which can often be exploited to run arbitrary code.

Rule	Severity	Likelihood	Remediation Cost	Priority	Level
ARR38-C	high	likely	medium	P18	L1

Related Vulnerabilities

The following vulnerability resulting from a violation of this rule is documented in the CERT Coordination Center Vulnerability Notes Database [CERT/CC VND].

Metric	ID	Date Public	Name
0	VU#162289	04/17/2006	C compilers may silently discard some wrap-around checks

References

- [Banahan 03] Section 5.3, “Pointers,” and Section 5.7, “Expressions Involving Pointers”
- [ISO/IEC 9899:1999] Section 6.5.6, “Additive Operators”
- [ISO/IEC PDTR 24772] “XYX Boundary Beginning Violation” and “XYZ Unchecked Array Indexing”
- [MITRE 07] CWE ID 129, “Unchecked Array Indexing”
- [VU#162289]

Risk Assessment

Failure to correctly determine the size of a wide character string can lead to buffer overflows and the execution of arbitrary code by an attacker.

Rule	Severity	Likelihood	Remediation Cost	Priority	Level
STR33-C	high	likely	medium	P18	L1

References

- [Viega 05] Section 5.2.15, “Improper String Length Checking”
- [ISO/IEC 9899:1999] Section 7.21, “String Handling <string.h>”
- [MITRE 07] CWE ID 119, “Failure to Constrain Operations within the Bounds of an Allocated Memory Buffer,” and CWE ID 135, “Incorrect Calculation of Multi-Byte String Length”
- [Seacord 05a] Chapter 2, “Strings”

■ STR34-C. Cast characters to unsigned types before converting to larger integer sizes

Signed character data must be converted to an unsigned type before being assigned or converted to a larger signed type. Because compilers have the latitude to define `char` to have the same range, representation, and behavior as either `signed char` or `unsigned char`, this rule should be applied to both `signed char` and (plain) `char` characters.

This rule is applicable only in cases where the character data may contain values that can be interpreted as negative values. For example, if the `char` type is represented by a two’s complement 8-bit value, any character value greater than +127 is interpreted as a negative value.

Noncompliant Code Example

This noncompliant code example is taken from a vulnerability in `bash` versions 1.14.6 and earlier that resulted in the release of CERT Advisory CA-1996-22. This vulnerability resulted from the sign extension of character data referenced by the `string` pointer in the `yy_string_get()` function in the `parse.y` module of the `bash` source code:

```
static int yy_string_get() {
    register char *string;
    register int c;

    string = bash_input.location.string;
    c = EOF;
```

```
/* If the string doesn't exist, or is empty, EOF found. */
if (string && *string) {
    c = *string++;
    bash_input.location.string = string;
}
return (c);
}
```

The `string` variable is used to traverse the character string containing the command line to be parsed. As characters are retrieved from this pointer, they are stored in a variable of type `int`. For compilers in which the `char` type defaults to `signed char`, this value is sign-extended when assigned to the `int` variable. For character code 255 decimal (−1 in two's complement form), this sign extension results in the value −1 being assigned to the integer, which is indistinguishable from EOF.

This problem was repaired by explicitly declaring the `string` variable as `unsigned char`.

```
static int yy_string_get() {
    register unsigned char *string;
    register int c;

    string = bash_input.location.string;
    c = EOF;

    /* If the string doesn't exist, or is empty, EOF found. */
    if (string && *string) {
        c = *string++;
        bash_input.location.string = string;
    }
    return (c);
}
```

This solution, however, is in violation of STR04-C, “Use plain `char` for characters in the basic character set.”

Compliant Solution

In this compliant solution, the result of the expression `*string++` is cast to (`unsigned char`) before assignment to the `int` variable `c`.

```
static int yy_string_get() {
    register char *string;
    register int c;

    string = bash_input.location.string;
    c = EOF;
```

```

/* If the string doesn't exist, or is empty, EOF found. */
if (string && *string) {
    /* cast to unsigned type */
    c = (unsigned char)*string++;

    bash_input.location.string = string;
}
return (c);
}

```

Risk Assessment

This is a subtle error that results in a disturbingly broad range of potentially severe vulnerabilities.

Rule	Severity	Likelihood	Remediation Cost	Priority	Level
STR34-C	medium	probable	medium	P8	L2

References

- [ISO/IEC 9899:1999] Section 6.2.5, “Types”
- [MISRA 04] Rule 6.1, “The plain char type shall be used only for the storage and use of character values”
- [MITRE 07] CWE ID 704, “Incorrect Type Conversion or Cast”

■ STR35-C. Do not copy data from an unbounded source to a fixed-length array

Functions that perform unbounded copies often rely on external input to be a reasonable size. Such assumptions may prove to be false, causing a buffer overflow to occur. For this reason, care must be taken when using functions that may perform unbounded copies.

Noncompliant Code Example (gets())

The `gets()` function is inherently unsafe and should never be used because it provides no way to control how much data is read into a buffer from `stdin`. This compliant code example assumes that `gets()` will not read more than `BUFSIZ - 1` characters from `stdin`. This is an invalid assumption, and the resulting operation can cause a buffer overflow.

According to Section 7.19.7.7 of C99, the `gets()` function reads characters from the `stdin` into a destination array until end-of-file is encountered or a new-line character is

References

- [ISO/IEC 9899:1999] Section 7.20.3.2, “The free Function”
- [ISO/IEC PDTR 24772] “DCM Dangling References to Stack Frames” and “XYK Dangling Reference to Heap”
- [Kernighan 88] Section 7.8.5, “Storage Management”
- [MISRA 04] Rule 17.6
- [MITRE 07] CWE ID 416, “Use After Free”
- [OWASP Freed Memory]
- [Seacord 05a] Chapter 4, “Dynamic Memory Management”
- [Viega 05] Section 5.2.19, “Using Freed Memory”

■ MEM31-C. Free dynamically allocated memory exactly once

Freeing memory multiple times has similar consequences to accessing memory after it is freed. The underlying data structures that manage the heap can become corrupted in a way that can introduce security vulnerabilities into a program. These types of issues are referred to as double-free vulnerabilities. In practice, double-free vulnerabilities can be exploited to execute arbitrary code. VU#623332, which describes a double-free vulnerability in the MIT Kerberos 5 function `krb5_recvauth()` [MIT 05] is one example.

To eliminate double-free vulnerabilities, it is necessary to guarantee that dynamic memory is freed exactly one time. Programmers should be wary when freeing memory in a loop or conditional statement; if coded incorrectly, these constructs can lead to double-free vulnerabilities. It is also a common error to misuse the `realloc()` function in a manner that results in double-free vulnerabilities (see MEM04-C, “Do not perform zero-length allocations”).

Noncompliant Code Example

In this noncompliant code example, the memory referred to by `x` may be freed twice: once if `error_condition` is true and again at the end of the code.

```
size_t num_elem = /* some initial value */;
int error_condition = 0;

int *x = (int *)malloc(num_elem * sizeof(int));
if (x == NULL) {
    /* handle allocation error */
}
/* ... */
if (error_condition == 1) {
    /* handle error condition*/
    free(x);
}
```

```

}
/* ... */
free(x);

```

Compliant Solution

In this compliant solution, the memory referenced by `x` is freed only once. This is accomplished by eliminating the call to `free()` when `error_condition` is equal to 1.

```

size_t num_elem = /* some initial value */;
int error_condition = 0;

if (num_elem > SIZE_MAX/sizeof(int)) {
    /* Handle overflow */
}
int *x = (int *)malloc(num_elem * sizeof(int));
if (x == NULL) {
    /* handle allocation error */
}
/* ... */
if (error_condition == 1) {
    /* Handle error condition */
}
/* ... */
free(x);
x = NULL;

```

Note that this solution checks for numeric overflow (see INT32-C, “Ensure that operations on signed integers do not result in overflow”).

Risk Assessment

Freeing memory multiple times can result in an attacker executing arbitrary code with the permissions of the vulnerable process.

Rule	Severity	Likelihood	Remediation Cost	Priority	Level
MEM31-C	high	probable	medium	P12	L1

References

- [ISO/IEC PDTR 24772] “XYK Dangling Reference to Heap” and “XYL Memory Leak”
- [MIT 05]
- [MITRE 07] CWE ID 415, “Double Free”

- [OWASP, Double Free]
- [Viega 05] “Doubly Freeing Memory”
- [VU#623332]

■ MEM32-C. Detect and handle memory allocation errors

The return values for memory allocation routines indicate the failure or success of the allocation. According to C99, `calloc()`, `malloc()`, and `realloc()` return null pointers if the requested memory allocation fails [ISO/IEC 9899:1999]. Failure to detect and properly handle memory management errors can lead to unpredictable and unintended program behavior. As a result, it is necessary to check the final status of memory management routines and handle errors appropriately.

Table 9–1 shows the possible outcomes of the standard memory allocation functions.

Noncompliant Code Example

In this noncompliant code example, `input_string` is copied into dynamically allocated memory referenced by `str`. However, the result of `malloc()` is not checked before `str` is referenced. Consequently, if `malloc()` fails, the program abnormally terminates.

```
char *input_string = /* initialize from untrusted data */;

size_t size = strlen(input_string) + 1;
char *str = (char *)malloc(size);
strcpy(str, input_string);
/* ... */
free(str);
str = NULL;
```

Table 9–1. Possible outcomes of standard memory allocation functions

Function	Successful Return	Error Return
<code>malloc()</code>	pointer to allocated space	null pointer
<code>calloc()</code>	pointer to allocated space	null pointer
<code>realloc()</code>	pointer to the new object	null pointer

■ FIO30-C. Exclude user input from format strings

Never call any formatted I/O function with a format string containing user input.

An attacker who can fully or partially control the contents of a format string can crash a vulnerable process, view the contents of the stack, view memory content, or write to an arbitrary memory location and consequently execute arbitrary code with the permissions of the vulnerable process [Seacord 05a].

Formatted output functions are particularly dangerous because many programmers are unaware of their capabilities (for example, they can write an integer value to a specified address using the `%n` conversion specifier).

Noncompliant Code Example

This noncompliant code example shows the `incorrect_password()` function, which is called during identification and authentication if the specified user is not found, or the password is incorrect, to display an error message. The function accepts the name of the user as a null-terminated byte string referenced by `user`. This is an excellent example of data that originates from an untrusted, unauthenticated user. The function constructs an error message, which is then output to `stderr` using the C99 standard `fprintf()` function [ISO/IEC 9899:1999].

```
#define MSG_FORMAT "%s cannot be authenticated.\n"
void incorrect_password(const char *user) {
    /* user names are restricted to 256 characters or less */
    static const char *msg_format = MSG_FORMAT;
    size_t len = strlen(user) + sizeof(MSG_FORMAT);
    char *msg = (char *)malloc(len);
    if (!msg) {
        /* Handle error condition */
    }
    int ret = snprintf(msg, len, msg_format, user);
    if (ret < 0 || ret >= len) {\
        /* Handle error */
    }
    fprintf(stderr, msg);
    free(msg);
    msg = NULL;
}
```

The `incorrect_password()` function calculates the size of the message, allocates dynamic storage, and constructs the message in the allocated memory using the `snprintf()` function. The addition operations are not checked for integer overflow because the length of the string referenced by `user` is known to have a length of 256 or less. Because the `%s` characters are replaced by the string referenced by `user` in the call to `snprintf()`, one less

byte is required to store the resulting string and terminating NULL-byte character. This is a common idiom for displaying the same message in multiple locations or when the message is difficult to build. The resulting code contains a format-string vulnerability, however, because the `msg` includes untrusted user input and is passed as the format-string argument in the call to `fprintf()`.

Compliant Solution (`fputs()`)

This compliant solution fixes the problem by replacing the `fprintf()` call with a call to `fputs()`, which does not treat `msg` like a format string but outputs it to `stderr` as is.

```
#define MSG_FORMAT "%s cannot be authenticated.\n"
void incorrect_password(const char *user) {
    /* user names are restricted to 256 characters or less */
    static const char *msg_format = MSG_FORMAT;
    size_t len = strlen(user) + sizeof(MSG_FORMAT);
    char *msg = (char *) malloc(len);
    if (!msg) {
        /* Handle error condition */
    }
    int ret = snprintf(msg, len, msg_format, user);
    if (ret < 0 || ret >= len) {\
        /* Handle error */
    }
    if (fputs(msg, stderr) == EOF) {
        /* Handle error */
    }
    free(msg);
    msg = NULL;
}
```

Compliant Solution (`fprintf()`)

This simpler compliant solution passes the untrusted user input as one of the variadic arguments to `fprintf()` and not as part of the format string, eliminating the possibility of a format-string vulnerability.

```
#define MSG_FORMAT "%s cannot be authenticated.\n"
void incorrect_password(char const *user) {
    fprintf(stderr, MSG_FORMAT user);
}
```

Noncompliant Code Example (POSIX)

This noncompliant code example is exactly the same as the first noncompliant code example but uses the POSIX function `syslog()` [Open Group 04] instead of the `fprintf()` function, which is also susceptible to format-string vulnerabilities.

```
#define MSG_FORMAT "%s cannot be authenticated.\n"
void incorrect_password(const char *user) {
    /* user names are restricted to 256 characters or less */
    static const char *msg_format = MSG_FORMAT;
    size_t len = strlen(user) + sizeof(MSG_FORMAT);
    char *msg = (char *)malloc(len);
    if (!msg) {
        /* Handle error condition */
    }
    int ret = snprintf(msg, len, msg_format, user);
    if (ret < 0 || ret >= len) {\
        /* Handle error */
    }
    syslog(LOG_INFO, msg);
    free(msg);
    msg = NULL;
}
```

The `syslog()` function first appeared in BSD 4.2 and is supported by Linux and other modern UNIX implementations. It is not available on Windows systems.

Compliant Solution (POSIX)

This compliant solution passes the untrusted user input as one of the variadic arguments to `syslog()` instead of including it in the format string.

```
#define MSG_FORMAT "%s cannot be authenticated.\n"
void incorrect_password(const char *user) {
    syslog(LOG_INFO, MSG_FORMAT user);
}
```

Risk Assessment

Failing to exclude user input from format specifiers may allow an attacker to crash a vulnerable process, view the contents of the stack, view memory content, or write to an arbitrary memory location, and consequently execute arbitrary code with the permissions of the vulnerable process.

Rule	Severity	Likelihood	Remediation Cost	Priority	Level
FIO30-C	high	likely	medium	P18	L1

Two recent examples of format-string vulnerabilities resulting from a violation of this rule include Ettercap (ettercap.sourceforge.net/history.php) and Samba (samba.org/samba/security/CVE-2007-0454.html). In Ettercap v.NG-0.7.2, the ncurses user interface suffers from a format string defect. The `curses_msg()` function in `ec_curses.c` calls `wdg_scroll_print()`, which takes a format string and its parameters and passes it to

`vw_printw()`. The `curses_msg()` function uses one of its parameters as the format string. This input can include user data, allowing for a format string vulnerability. The Samba AFS ACL mapping VFS plug-in fails to properly sanitize user-controlled file names that are used in a format specifier supplied to `snprintf()`. This security flaw becomes exploitable when a user can write to a share that uses Samba's `afsac1.so` library for setting Windows NT access control lists on files residing on an AFS file system.

Related Vulnerabilities. The following vulnerabilities resulting from the violation of this rule are documented in the CERT Coordination Center Vulnerability Notes Database [CERT/CC VND].

Metric	ID	Date Public	Name
1.81	VU#649732	02/05/2007	Samba AFS ACL mapping VFS plug-in format string vulnerability
11.85	VU#794752	01/20/2007	Apple iChat AIM URI handler format string vulnerability
1.8	VU#512491	03/05/2008	GNOME Evolution format string vulnerability
8.11	VU#286468	05/31/2005	Ettercap contains a format string error in the "curses_msg()" function

References

- [ISO/IEC 9899:1999] Section 7.19.6, “Formatted Input/Output Functions”
- [ISO/IEC PDTR 24772] “RST Injection”
- [MITRE 07] CWE ID 134, “Uncontrolled Format String”
- [Open Group 04] `syslog()`
- [Seacord 05a] Chapter 6, “Formatted Output”
- [Viega 05] Section 5.2.23, “Format String Problem”

■ FIO31-C. Do not simultaneously open the same file multiple times

Simultaneously opening a file multiple times has implementation-defined behavior. While some platforms may forbid a file simultaneously being opened multiple times, platforms that allow it may facilitate dangerous race conditions.

Noncompliant Code Example

This noncompliant code example logs the program's state at runtime.

Compliant Solution

This compliant solution implements a strictly conforming test for unsigned overflow.

```
unsigned int ui1, ui2, sum;

if (UINT_MAX - ui1 < ui2) {
    /* Handle error condition */
}
sum = ui1 + ui2;
```

If the noncompliant form of this test is truly faster, talk to your compiler vendor, because if these tests are equivalent, optimization should occur. If both forms have the same performance, prefer the portable form.

Risk Assessment

Unnecessary platform dependencies are, by definition, unnecessary. Avoiding these dependencies can eliminate porting errors resulting from invalidated assumptions.

Recommendation	Severity	Likelihood	Remediation Cost	Priority	Level
MSC14-C	low	unlikely	medium	P2	L3

References

- [Dowd 06] Chapter 6, “C Language Issues” (Arithmetic Boundary Conditions, pp. 211–223)
- [ISO/IEC 9899:1999] Section 3.4.1, “Implementation-Defined Behavior,” Section 3.4.4, “Unspecified Behavior,” Annex J.1, “Unspecified Behavior,” and Annex J.3, “Implementation-Defined Behavior”
- [ISO/IEC PDTR 24772] “BQF Unspecified Behaviour”
- [Seacord 05a] Chapter 5, “Integers”

■ MSC15-C. Do not depend on undefined behavior

C99, Section 3.4.3, defines *undefined behavior* as

behavior, upon use of a nonportable or erroneous program construct or of erroneous data, for which this International Standard imposes no requirements

C99, Section 4, explains how the standard identifies undefined behaviors:

If a “shall” or “shall not” requirement that appears outside of a constraint is violated, the behavior is undefined. Undefined behavior is otherwise indicated in this International Standard by the words “undefined behavior” or by the omission of any explicit definition of behavior. There is no difference in emphasis among these three; they all describe “behavior that is undefined.”

C99, Annex J.2, “Undefined Behavior,” contains a list of explicit undefined behaviors in C99.

Behavior can be classified as undefined by the C standards committee for the following reasons:

- to give the implementor license not to catch certain program errors that are difficult to diagnose.
- to identify areas of possible conforming language extension: the implementor may augment the language by providing a definition of the officially undefined behavior.

Conforming implementations can deal with undefined behavior in a variety of fashions, such as ignoring the situation completely, with unpredictable results; translating or executing the program in a documented manner characteristic of the environment (with or without the issuance of a diagnostic message); or terminating a translation or execution (with the issuance of a diagnostic message). Because compilers are not obligated to generate code for undefined behavior, these behaviors are candidates for optimization. By assuming that undefined behaviors will not occur, compilers can generate code with better performance characteristics.

Unfortunately, undefined behaviors are coded. Optimizations make it difficult to determine how these systems will behave in the presence of undefined behaviors. This is particularly true when visually inspecting source code that relies on undefined behaviors; a code reviewer cannot be certain if the code will be compiled or if it will be optimized out. Furthermore, just because a compiler currently generates object code for an undefined behavior does not mean that future versions of the compiler are obligated to do the same; the behavior may be viewed as an opportunity for further optimization. Compilers are also not required to issue diagnostics for undefined behavior, so there is frequently no easy way to identify undefined behavior in code.

All of this puts the onus on the programmer to write strictly conforming code, with or without the help of the compiler. Because performance is a primary emphasis of the C language, this situation is likely to get worse before it gets better.

Noncompliant Code Example

An example of undefined behavior in C99 is the behavior on signed integer overflow. This noncompliant code example depends on this behavior to catch the overflow.

```
#include <assert.h>

int foo(int a) {
    assert(a + 100 > a);
    printf("%d %d\n", a + 100, a);
    return a;
}

int main(void) {
    foo(100);
    foo(INT_MAX);
}
```

This code tests for signed integer overflow by testing to see if $a + 100 > a$. This test cannot evaluate to false unless an integer overflow occurs. However, because a conforming implementation is not required to generate code for undefined behavior, and signed integer overflow is undefined behavior, this code may be compiled out. For example, GCC 4.1.1 optimizes out the assertion for all optimization levels, and GCC 4.2.3 optimizes out the assertion for programs compiled with `-O2`-level optimization and higher.

On some platforms, the integer overflow will cause the program to terminate (before it has an opportunity to test).

Compliant Solution

This compliant solution does not depend on undefined behavior.

```
#include <assert.h>

int foo(int a) {
    assert(a < (INT_MAX - 100));
    printf("%d %d\n", a + 100, a);
    return a;
}

int main(void) {
    foo(100);
    foo(INT_MAX);
}
```

Risk Assessment

While it is rare that the entire application can be strictly conforming, the goal should be that almost all the code is allowed for a strictly conforming program (which among other things means that it avoids undefined behavior), with the implementation-dependent parts confined to modules that the programmer knows he or she needs to adapt to the platform when it changes.

Recommendation	Severity	Likelihood	Remediation Cost	Priority	Level
MSC15-C	high	likely	medium	P18	L1

Related Vulnerabilities. The following vulnerability resulting from the violation of this recommendation is documented in the CERT Coordination Center Vulnerability Notes Database [CERT/CC VND].

Metric	ID	Date Public	Name
0	VU#162289	04/17/2006	C compilers may silently discard some wraparound checks

References

- [ISO/IEC 9899:1999] Section 3.4.3, “Undefined Behavior,” Section 4, “Conformance,” and Annex J.2, “Undefined Behavior”
- [ISO/IEC PDTR 24772] “BQF Unspecified Behaviour,” “EWF Undefined Behaviour,” and “FAB Implementation-Defined Behaviour”
- [Seacord 05a] Chapter 5, “Integers”

■ MSC30-C. Do not use the rand() function for generating pseudorandom numbers

Pseudorandom number generators use mathematical algorithms to produce a sequence of numbers with good statistical properties, but the numbers produced are not genuinely random.

The C Standard function rand() (available in `stdlib.h`) does not have good random number properties. The numbers generated by rand() have a comparatively short cycle, and the numbers may be predictable.

Index

- , (commas), surrounding macro names, 12
- . (period), unique header file names, 24
- /*...*/, comment delimiters, 578–580
- & (ampersand), bitwise AND operator, 174–175
- && (ampersands), logical AND, 96–98, 113
- \ (backslash), loop terminator, 260
- ^ (caret), bitwise OR operator, 174–175
- (dash), in file names, 591
- = (equal sign), equality operator, 283
- == (equal signs), equality operator, 283
- != (exclamation, equal sign), equality operator, 283
- >> (greater-than signs), right-shift operator, 174–175
- << (less-than signs), left-shift operator, 174–175
- % (percent sign), remainder operator, 168–170
- ? (question mark), repeating, 22–24
- | (vertical bar), bitwise inclusive OR operator, 174–175
- || (vertical bars), logical OR, 96–98, 113
- [] (square brackets), subscript operator, 242
- # (pound sign) operator, concatenating tokens, 12, 18–20
- ## (pound signs) operator, stringifying macro parameters, 12, 18–20
- ~ (tilde), complement operator, 174–175
- 3Rs of survivability, 534
- A**
- abort()
 - called by signal handlers, 564
 - exit behavior, 547, 556–557
 - terminating a program, 546–549
- Abstract data types. *See* Data types, abstract.
- Abstract machine, 582
- Abstraction levels, memory allocation, 319–321
- Access control
 - files. *See* Files, access control.
 - objects, 82–84
 - permission escalation
 - signal handlers, 516
 - temporary file creation, 463
 - TOCTOU race conditions, 635
- Access control, permissions
 - appropriate level of, 394–397
 - improper, information leakage, 417
 - principle of least privilege, 620–623
 - relinquishing, 636–642
- Access-freed-memory vulnerability, 322–323
- Accessing files
 - with file descriptors, 372–374
 - by file name, 372–374
 - with FILE pointers, 372–374
 - unintentionally, 374, 383–386, 463
 - objects outside their lifetime, 72–74
- acos(), 229
- acosh(), 229
- Actual argument, 19
- Addition
 - integer overflow, 192–194
 - unsigned integer wrapping, 182–183
- Address arguments, 552, 555
- Adobe Flash vulnerability, 357
- Advisory file locks, 456

- Alignment
 - bit-fields, 115–116
 - members, 98–100
 - objects, 131–133
 - struct member, 98–100
 - structures in memory, 100–102
 - type alignment error, 63
- Allocating memory. *See* Memory, allocating.
- Ampersand (&), bitwise AND operator, 174–175
- Ampersands (&&), logical AND, 96–98, 113
- AND operators
 - & (ampersand), bitwise AND, 174–175
 - && (ampersands), logical AND, 96–98, 113
- ANSI/IEEE 754-2007, 215
- arc4random(), 609
- Arguments
 - assignment, 30–32
 - decrement, 30–32
 - function calls, 113, 133–135
 - increment, 30–32
 - passing with function pointers, 133–134
 - variadic functions, 60–61
 - volatile access, 30–32
- argv(), 295–298
- argv[0], 295–298
- Ariane 5 launcher failure, 164
- Arithmetic conversions, 150–151
- Arithmetic operations on array pointers, 261–263, 263–269
- Arrays
 - \ (backslash), loop terminator, 260
 - [] (square brackets), subscript operator, 242
 - adding pointers, 263–269, 265–269
 - bounds
 - checking, 256
 - specifying, 248–250, 312–313
 - buffer overflows
 - copying into insufficient storage, 255–258
 - incorrect use of sizeof(), 247
 - looping beyond last element, 259–261
 - characters and strings
 - bounds specification, 312–313
 - initialization, 312–313
 - unbounded copying, 307–311
 - comparing pointers, 261–263
 - copying data into, 255–258
 - definition, 243
 - description, 242–245
 - dynamic allocation
 - for copied data, 256–257
 - resizing, 343–346
 - size requirements, calculating, 342–343
 - element types, 243, 258
 - elements
 - definition, 243
 - dereferencing, 264
 - finding, 259–261
 - loop termination, 259–261
 - flexible members, 358–360
 - incompatible types, 258–259
 - incomplete types, 243
 - indices, range checking, 250–251
 - initializing
 - array bounds, specifying, 248–250
 - incomplete types, 243
 - string literals, 312–313
 - iterating through, 259–261
 - memory overwrites, 251, 253
 - modification failure, resetting strings, 446–447
 - notation consistency, 251–254
 - null-terminated byte strings, 259–261
 - pointers
 - adding, 263–269, 265–269
 - arithmetic operations on, 261–263, 263–265, 265–269
 - comparing, 261–263
 - subtracting, 261–263, 263–265, 265–269
 - range checking
 - indices, 250–251
 - size arguments, 254–255
 - risk assessment summary, 242
 - rules and recommendations
 - related, 242
 - summary of, 241
 - size
 - arguments, variable length arrays, 254–255
 - bounds specification, 312–313
 - determining, 95–96, 245–249
 - insufficient, 255–258
 - string literal initialization, 312–313
 - stack exhaustion, 255
 - subscripts, and pointers, 244
 - subtracting pointers, 261–263, 263–265, 265–269
 - unbounded copying, 307–311
 - variable length
 - size arguments, 254–255
 - stack allocation, 335–336
- ASCII characters, 590–593
- asin(), 229
- assert()
 - diagnostic tests, 597–598
 - invalid pointers, terminating on, 350
 - side effects, 122–123
 - termination behavior, 556–557
 - testing values, constant expressions, 39–42
- Assertions
 - side effects, 122–123
 - static. *See* Static assertions.
 - termination behavior, 556–557
 - testing code, 597–598
- Assignment side effects, 30–32
- Asynchronous-signal-safe functions, 511–516. *See also* Longjmp(), in signal handlers.
- atan2(), 229
- atanh(), 229
- atexit(), 494–497, 545
- atof(), 559
- atoi(), 160–162
- atol(), 160–162
- ato11(), 160–162
- Automatic storage duration
 - buffers, pointers to, 631–633
 - const-qualified objects, 46
 - object initialization, 124
 - object lifetime, 72
- Automatically generated code, 2–3
- Availability, measuring, 533
- B**
- Backslash (\), loop terminator, 260
- bash vulnerability, 305
- Basic character set, 29, 282–284, 590
- Binary data
 - floating-point, 215
 - reading, 401–403
 - transferring across systems, 402–403
 - writing, 401–403
- Binary mode vs. text mode, file streams, 411–413, 442–443

- Binary temporary files, creating, 460–461
- Bit-fields
 - alignment, 115–116
 - layout assumptions, 172–173
 - in multithreaded environments, 626–629
 - overlap, 116–117
 - type assumptions, 172–173
- Bitwise operations
 - & (ampersand), bitwise AND, 174–175
 - ^ (caret), bitwise OR, 174–175
 - | (vertical bar), bitwise inclusive OR, 174–175
 - combined with arithmetic operations, 175–178
 - on unsigned operands, 174–175
- Black listing, characters and strings, 278
- Block scope, 73
- Books and publications. *See* References.
- Bounds. *See also* Range.
 - checking
 - arrays, 256
 - `calloc()` arguments, 342–343
 - floating-point math functions, 229
 - memory allocation, 342–343
 - specifying, arrays, 248–250, 312–313
- Buffer overflow causes
 - arrays
 - copying into insufficient storage, 255–258
 - incorrect use of `sizeof()`, 247
 - looping beyond last element, 259–261
 - characters and strings
 - alternative functions, 289–291
 - null-terminated byte strings, 280–282, 294–299
 - unbounded copying, 307–311
 - wide string size, 303–305
 - data model errors, 145
 - error handling, return codes, 436
 - insufficient storage for variables, 474
 - invalid size arguments, 365
 - magic numbers, 50
 - memory allocation, 343
 - memory initialization, 346–348
 - numeric literals, 50
 - pointer arithmetic, 107–109
 - programmer-defined integer types, 180
 - `readlink()` termination, 624
 - return code checking, 436
 - truncation errors, 190
 - unsigned integer wrapping, 343
- `BUFSIZ` macro, 106
- Building Systems from Commercial Components*, xxxi
- C**
 - Caching restrictions, 82–84
 - Call-by-name, 11
 - Call-by-value, 11
 - `calloc()`
 - arguments, out-of-bounds check, 342–343
 - clearing freed space, 330
 - information leakage, 330
 - memory allocation errors, 355–357
 - memory size, calculating, 362–365
 - returned pointers, casting, 324–328
 - zero-length allocations, 332–333
 - Canonical form conversion, path names
 - Linux, 380–381
 - overview, 374–375
 - POSIX, 376–380
 - Windows, 381–382
 - `canonicalize_file_name()`, 380–381
 - Caret (^), bitwise OR operator, 174–175
 - Carriage returns, 412, 591
 - Case label, 574
 - Case sensitivity
 - environment variables, 475, 477
 - file names, 591
 - header file names, 24
 - Casting. *See also* Conversions.
 - `calloc()` returned pointers, 324–328
 - function calls into pointers, 10, 324–328
 - memory allocation calls to pointers, 10
 - `realloc()` returned pointers, 324–328
 - Casting away
 - `const`-qualification, 35, 102–104
 - volatile qualifications, 123–124
 - CERT Secure Coding Standards
 - wiki, xvii
 - CERT/CC VND (CERT/CC Vulnerability Notes Database), xxiv
 - char type, 162–164
 - Characters and strings
 - arrays
 - bounds specification, 312–313
 - initialization, 312–313
 - unbounded copying, 307–311
 - ASCII characters, 590–593
 - basic character set, 29, 282–284, 590
 - black listing, 278
 - buffer overflows
 - alternative functions, 289–291
 - null-terminated byte strings, 280–282, 294–299
 - unbounded copying, 307–311
 - wide string size, 303–305
 - C language support for, 273
 - character class ranges, 370–371
 - character encoding
 - ASCII characters, 590–593
 - Unicode, 594–597
 - UTF-8, 594–597
 - character I/O functions, return code checking, 436–438
 - character literals, 45
 - character sets
 - basic, 29, 282–284, 590
 - execution, 590
 - source, 590
 - character-handling functions, 314–315
 - comparing, 274
 - in complex subsystems, 276–280
 - converting to integers, 143, 157–162. *See also* Stringification.
 - data integrity issues, 279
 - data sanitization, 276–280, 292, 590–593
 - data types. *See also specific types.*
 - character types, 162–164, 282–284
 - choosing, 273–275
 - converting, 305–307
 - int, 274
 - for numeric values, 162–164, 282–284
 - plain char, 162–164, 274, 282–284

- Characters and strings, *continued*
 - data types, *continued*
 - signed character, 162–164, 274, 305–307
 - unsigned character, 162–164, 274, 305–307
 - wchar_t, 274
 - execution character set, 590
 - extended characters, 590
 - function alternatives, 288–291
 - functions for, arguments to, 314–315
 - injection attacks, 279
 - length, 273
 - loss of data, 279
 - managed strings, 291–293
 - managing consistently, 275–276
 - memory usage, reducing, 301–303
 - multibyte strings, 273
 - narrow strings, 273, 284
 - new manipulation code, 291–293
 - null-terminated byte strings
 - buffer overflows, 280–282, 294–299
 - creating, 299–303
 - definition, 273
 - dynamic allocation, 275–276
 - environment variables, copying, 298–299
 - libraries, 275–276
 - managing consistently, 275–276
 - null termination character, 299–303
 - off-by-one error, 294–295
 - static allocation, 275–276
 - storage allocation, 294–299
 - truncating, 280–282, 300–301
 - unintended character arrays, 312–313
 - pointers to, 273
 - risk assessment summary, 272
 - rules and recommendations related, 273
 - summary of, 271–272
 - single-byte strings, 273
 - source character set, 590
 - special characters
 - ASCII, 590
 - national use positions, 590
 - portability, 590
 - sanitizing, 276–280
 - string data, converting to integers, 143, 157–162
 - string tokenization, 286–287
 - unbounded copying, 307–311
 - US-ASCII characters, 590–593
 - white listing, 277–278
 - wide strings
 - data type, 274
 - definition, 273
 - modifying, 293–294
 - sizing, 303–305
 - string literals, 285
- Characters and strings, string literals
 - arrays, initializing, 312–313
 - concatenating, 12
 - const-qualification, 67, 284–285
 - definition, 45, 293
 - immutable, 284–285
 - modifying, 67, 293–294
 - mutable, 284–285
 - narrow, 284
 - size specification, 312–313
 - wide
 - data type, 274
 - definition, 273
 - modifying, 293–294
 - sizing, 303–305
 - string literals, 285
- chroot() jails, 419–420
- Clare, Geoff, 482
- Clean compiles, 570–572
- clearenv(), 479–481
- Clearing
 - the environment, 478–482
 - returned resources, 328–332, 582–585
- Close-on-exec flag, 452
- Closing files, 450–454
- Code
 - /*...*/, comment delimiters, 578–580
 - analysis tools, xx
 - assertions, 597–598
 - categories, 2–3
 - commenting out, 578–579
 - comments, 578–580
 - compiling. *See* Compiling code.
 - errors of addition, 576–578
 - errors of omission, 574–576
 - logical completeness, 572–574
 - platform dependencies, avoiding, 602–604
 - portability, 602–604
 - readability, 579–580
 - testing
 - with assertions, 597–598
 - compiler diagnostic messages, 571
 - constant expression values, 19, 39–42
 - functions for error conditions, 57–59
 - undefined behavior, 604–607
 - unnecessary, removing, 582–588, 598–600
 - unused values, removing, 600–602
- Coding standards
 - automatically generated code, 2–3
 - code categories, 2–3
 - deviation procedures, 4
 - hand-coded code, 2–3
 - rules and recommendations, levels, 4
 - source code, compliance validation, 3–4
 - system qualities, 1–2
 - tool selection, 3–4
 - tool-generated code, 2–3
- Coding style, xxii
- Command processors, 482–487
- Commas (,), surrounding macro names, 12
- Commenting out code, 578–579
- Comments
 - /*...*/, comment delimiters, 578–580
 - commenting out code, 578–579
 - describing code, 578–580
 - readability, 42–43, 579–580
 - variable declarations, 42–43
- Common Vulnerabilities and Exposures (CVE), xxiv
- Common Vulnerability Scoring System (CVSS), xxvi
- Common Weakness Enumeration (CWE) IDs, xxiv
- Comparing
 - array pointers, 261–263
 - characters and strings, 274
 - integers, to a larger size, 207–210
 - mbstowcs() return values, 611
 - return values, 610–613
 - size_t return values, 611
 - structures, 100–102
- Compile-time constant, yielding, 10

- Compiling code
 - clean compiles, 570–572
 - diagnostic messages, 571
 - optimization, and sensitive information, 582–585
 - warning levels, 570–572
- `complex.h`, 226, 559
- Compliance with standards, 1, 3–4
- Compliant solutions, xxiii
- Concatenation
 - # (pound sign) operator, concatenating tokens, 12, 18–20
 - string literals, 12
 - token pasting, 19
 - tokens, 12, 18–20, 29–30
 - universal character names, 29–30
- Conditional inclusion, 41
- Conditional operator, 151
- `confstr()`, environment variables, 479, 481
- `const` poisoning, 35
- Constant values
 - assuming, 105–106
 - modifying, 102–104
 - testing, 19, 39–42
- Constants, 45. *See also* Characters and strings, string literals; Declarations and initialization, literals.
- `const-qualification`
 - casting away, 35, 102–104
 - function parameters, 66–68
 - immutable objects, 14, 35–36
 - objects, scope, 45–46
 - scope, 42
 - string literals, 284–285
- Constraint, 541–543
- Control flow, altering on error detection, 554–555
- Conversion specifiers
 - argument mismatches, 371
 - length modifiers, 370
 - list of, 370
- Conversions. *See also* Casting; Integer type conversions.
 - arithmetic, 150–151
 - characters to integers, 143, 157–162. *See also* Stringification.
 - data types, characters and strings, 305–307
 - data types, integer
 - arithmetic conversions, 150–151
 - code examples, 151–153
 - data loss, 186–191
 - integer conversion rank, 150
 - integer promotions, 149–150, 152
 - integer-to-pointer, 170–172
 - loss of precision, 188–189
 - minimum ranges, 189–190
 - misinterpreted data, 186–191
 - numeric strings, to greatest-width integers, 143
 - pointer-to-integer, 170–172
 - required conversions, determining, 150
 - rules for, 149–152
 - signed to signed, 188–189
 - signed to unsigned, 188
 - string data, 143, 157–162
 - unsigned to signed, 187
 - unsigned to unsigned, 188–189
 - usual arithmetic conversions, 150–151
 - directory names to canonical forms. *See* Canonical form conversion, path names.
 - expression pointers, 131–133
 - file names to canonical forms. *See* Canonical form conversion, path names.
 - floating-point errors, 219
 - floating-point numbers to integers, 234–239
 - function pointer types, 84–86
 - implicit, 324
 - integer conversion rank, 150
 - integers and floating-point numbers, 219
 - integer-to-pointer, 170–172
 - loss of precision, 188–189
 - memory pointers, unintended, 328
 - numeric, 559
 - numeric strings to greatest-width integers, 143
 - pointer-to-integer, 170–172
 - signed integers, 188–189
 - strings to integers, 143, 157–162. *See also* Stringification.
 - unsigned integers, 187–189
 - usual arithmetic, 150–151
- Copying
 - characters and strings, unbounded, 307–311
 - data into arrays, 255–258
 - environment variables, 298–299, 471–473
 - FILE objects, 443–444
 - null-terminated byte strings, 296–297
 - overlapping objects, 81
- Core dumps, 338–339
- `cosh()`, 229
- Creating
 - file names
 - hard coded, 457–458
 - predictability, 454–455
 - uniqueness, 454–455, 458–460, 459–460, 461–462
 - format strings, 370–371
 - null-terminated byte strings, 299–303
 - universal character names, 29–30
- Creating, files. *See also* `Fopen()`.
 - access permissions, 394–397
 - assumptions about, 383–387
 - hard-coded names, 457–459
 - mode strings, 407–408
 - in shared directories, 457–459
 - temporary
 - binary, 460–461
 - with hard-coded file name, 457–459
 - in shared directories, 454–463
 - unique file names, 454–455, 458–460, 459–460, 461–462
- `CryptGenRandom()`, 609
- `ctype.h`, 274, 314
- CVE (Common Vulnerabilities and Exposures), xxiv
- CVSS (Common Vulnerability Scoring System), xxvi
- CWE (Common Weakness Enumeration) IDs, xxiv
- D**
- Dangling pointers
 - accessing freed memory, 351–353
 - definition, 351
 - `free()`, 322–323
 - `getenv()`, 468
- Dash (-), in file names, 591
- Data integrity, characters and strings, 279
- Data loss
 - characters and strings, 279
 - information leakage
 - file descriptor leakage, 450–454
 - from freed resources, 328–332

- Data loss, *continued*
 - information leakage, *continued*
 - improper permissions, 417
 - multithreaded environments, 629
 - uninitialized memory, 346–348
 - overwriting data, 473
 - pushing back characters, 411
 - sensitive information
 - clearing, 328–332
 - compiler optimization, 582–585
 - improper permission vulnerability, 417
 - leaking, 328–332, 346–348, 417
 - program termination vulnerability, 549, 557
 - writing to disk, 338–341
- truncation
 - buffer overflow cause, 190
 - `fgetc()`, 280–282
 - `gets()`, 280–282
 - improper use of `strtok()`, 286–287
 - new-line character, missing, 440–442
 - null-terminated byte strings, 280–282, 300–301
 - pushing back characters, 411
 - `snprintf()`, 280–282
 - `sprintf()`, 280–282
 - `strcat()`, 280–282
 - `strcpy()`, 280–282
 - `strcpy_s()`, 281–282
 - string tokenization, 286–287
 - `strncat()`, 280–282
 - `strncpy()`, 280–282
 - toward zero, 169
 - `ungetc()`, 411
 - `ungetwc()`, 411
- Data models, integers
 - code examples, 143–145
 - definition, 141
 - integral ranges, 142
 - list of, 142
 - numeric strings, converting to greatest-width integers, 143
 - size restrictions, 142–143
 - types available, 142–143
- Data transfer across systems, 402–403
- Data types. *See also specific types.*
 - abstract, 64–65
 - alignment error, variadic functions, 63
 - arrays
 - incompatible types, 258–259
 - incomplete types, 243
 - bit-field alignment, 115–116
 - bit-field overlap, 116–117
 - characters and strings. *See also specific types.*
 - character types, 162–164, 282–284
 - choosing, 273–275
 - converting, 305–307
 - `int`, 274
 - for numeric values, 162–164, 282–284
 - plain char, 162–164, 274, 282–284
 - signed character, 162–164, 274, 305–307
 - unsigned character, 162–164, 274, 305–307
 - `wchar_t`, 274
 - choosing, 273–275
 - converting, 305–307
 - definitions
 - formatted I/O, 178–181
 - vs. macro definitions, 15–16
 - readability, 44–45
 - scope rules, 15
 - encoding, macro definitions vs. type definitions, 15–16
 - function pointers, 53–54
 - incompatibilities, 114–117
 - information hiding, 64–65
 - integer, conversions
 - arithmetic conversions, 150–151
 - code examples, 151–153
 - data loss, 186–191
 - integer conversion rank, 150
 - integer promotions, 149–150, 152
 - integer-to-pointer, 170–172
 - loss of precision, 188–189
 - minimum ranges, 189–190
 - misinterpreted data, 186–191
 - numeric strings, to greatest-width integers, 143
 - pointer-to-integer, 170–172
 - required conversions, determining, 150
 - rules for, 149–152
 - signed to signed, 188–189
 - signed to unsigned, 188
 - string data, 143, 157–162
 - unsigned to signed, 187
 - unsigned to unsigned, 188–189
 - usual arithmetic conversions, 150–151
 - integers vs. floating-point, 114–115
 - interpretation error, 62–63
 - for numeric values, 162–164, 282–284
 - opaque, 64–65
 - plain char, 162–164, 274, 282–284
 - size
 - determining, 49, 95–96, 109–111
 - hard coding, 109–111
 - variadic functions, 62–64
 - volatile qualifications, casting away, 123–124
- Date public, vulnerabilities, xxvii
- Decimal constant, 214
- Decimal floating-point, 215
- Declarations and initialization
 - abstract data types, 64–65
 - `assert()`, 39–42
 - casting away `const`-qualification, 35
 - comments, 42–43
 - `const` poisoning, 35
 - constant expression values, testing, 19, 39–42
 - `const`-qualification
 - casting away, 35
 - function parameters, 66–68
 - immutable objects, 14, 35–36
 - function declarators, type information, 51–54
 - function parameters, 66–68
 - function pointers
 - referencing overlapping objects, 80–82
 - `restrict` qualification, 80–82
 - type conversion, 84–86
 - type information, declaring, 53–54
- function prototypes, 52–53
- functions
 - `errno` error codes, 57–59
 - `errno_t` error code, 57–59
 - testing for error conditions, 57–59

- functions, variadic
 - type alignment error, 63
 - type interpretation error, 62–63
 - type issues, 62–64
 - va_arg() macro, 61, 62–64
 - vfprintf(), 61
 - vfscanf(), 61
 - vscanf(), 61
 - vsprintf(), 61
 - vsscanf(), 61
 - writer/caller contract, 59–62
- identifiers
 - conflicting linkage classifications, 87–89
 - duplicates, 78
 - external linkage, 87–89
 - implicit declaration, 4–77
 - internal linkage, 87–89
 - no linkage, 87–89
 - significant characters, minimum requirements, 78
 - uniqueness, 78–79
 - visual distinction, 38–39
- immutable objects, const-qualifying, 14, 35–36
- incorrect assumptions, identifying, 39–42
- information hiding, 64–65
- literals
 - character, 45
 - constants, 45
 - const-qualified objects, 45–46
 - enumeration, 46, 48–49
 - floating, 45
 - integer, 45, 48–49
 - kinds of, 36
 - magic numbers, 45
 - meaningful names, 45–51
 - memory requirements, 46–47
 - named, symbolic, 45–47, 50
 - object-like macros, 47
 - relationships, encoding, 54–56
 - runtime modification, 36
 - string, 45
- non-prototype-format declarators, 51–52
- objects
 - access restrictions, 82–84
 - accessing outside their lifetime, 72–74
 - caching restrictions, 82–84
 - confined to current scope, 70–72
 - declaring as static, 70–72
 - hiding, 70–72
 - lifetime duration, 72–74
 - lifetime-qualified type, 82–84
 - opaque data types, 64–65
 - readability
 - comments, 42–43
 - literals, 45
 - magic numbers, 45
 - meaningful literal names, 45–51
 - type definitions, 44–45
 - typedef, 44–45
 - variables, one per declaration, 42–44
 - risk assessment summary, 34–35
 - rules and recommendations, summary of, 33–34
 - runtime error checking, 39
 - sizeof operator, 49
 - static assertions, testing constant values, 19, 39–42
 - static_assert() macro, 41–42
 - translation units, global variable initialization order, 69–70
 - type size, determining, 49
 - variable names in subscopes, reusing, 8, 36–38
- variables
 - global initialization order, assumptions about, 69–70
 - multiple per declaration, 16, 42–44
 - reusing names in subscopes, 36–38
 - size, determining, 49
 - unintentional reference, 36–39
 - visually indistinct identifiers, 38–39
- Decrement side effects, 30–32
- Default argument promotions, 62
- Default label, 574
- Default values, looking up, 479
- #define vs. typedef, 15–16
- Deleting. *See* Removing.
- Delivering Signals for Fun and Profit*, 519
- Demotion of floating-point numbers, 236–239
- Denial-of-service attacks
 - code readability, 580
 - divide-by-zero errors, 156
 - FILE object copies, 444
 - file operations, on devices, 426
 - input/output, 426
 - memory allocation failures, 355–357
 - memory management errors, 321
 - race conditions, multiple threads, 615–616
 - reading freed memory, 352
 - stack allocation, 337–338
 - UTF8-encoded data, 596
 - vfork(), 629
- Dependability, measuring, 533
- Dereferencing array elements, 264
- Development process for secure software, xvi–xvii
- Deviation procedures, for rules and recommendations, 4
- Devices, accessing as files, 426–431
- Diagnostic messages, compiler, 571
- difftime(), 581
- Digraphs, 23
- Directories
 - names, converting to canonical forms. *See* Canonical form conversion, path names.
 - secure, 413–418
 - shared
 - creating files in, 457–459
 - temporary files in, 454–463
- Disk files, identifying, 430
- Divide-by-zero errors
 - denial of service attacks, 156
 - division operations, 201
 - floating-point numbers, 219, 228
 - modulo operations, 201–203
- Division
 - fractional remainders. *See* Modulo operations.
 - integer overflow, 197
 - modulo arithmetic, 165
 - modulo operations
 - divide-by-zero errors, 201–203
 - division, 169
 - fractional remainders, discarding, 169
 - integer overflow, 197–198
 - by powers of 2, 176–177
 - remainder assumptions, 168–170
 - with shift operators, 176–177
 - truncation toward zero, 169
- Domain errors, floating-point numbers, 227–234

- Double-free vulnerability, 319–321, 322–323, 353–355
- do-while loops, 27–29
- _dupenv_s(), 470–471
- Duplicate environment variable names, 475–478
- Duplicate identifiers, 78
- Dynamic allocation, arrays
 - for copied data, 256–257
 - resizing, 343–346
 - size requirements, calculating, 342–343
- E**
- EDOM, 559
- EILSEQ, 560
- Element types, 243, 258
- Elements of arrays
 - definition, 243
 - dereferencing, 264
 - finding, 259–261
 - loop termination, 259–261
- Encoding errors, 434, 537, 553, 560
- Encoding types, macro definitions
 - vs. type definitions, 15–16
- Endianess, 401
- End-of-file
 - detecting, 436–438, 438–440
 - error indicators, 438–440
 - indicator, 436, 438
- enum type, 48, 572–573, 574
- Enumeration constants, mapping, 167–168
- Enumeration literals, 46, 48–49
- environ argument, 491
- Environment
 - buffer overflows, insufficient storage, 474
 - clearing, 478–482
 - command processors, 482–487
 - external programs, invoking, 478–482
 - overwriting data, 473
 - program termination, 494–497
 - risk assessment summary, 467–468
 - rules and recommendations
 - related, 468
 - summary of, 467
 - sanitizing, 478–482
 - system() calls, 482–487
- Environment list, 298, 475, 479
- Environment variables
 - ~ vs. ~username, 482
 - case sensitivity, 475, 477
 - copying, 471–473
 - default values, determining, 479
 - default values, looking up, 479
 - determining need for, 481
 - duplicate names, 475–478
 - getting, 468–471
 - \$HOME value, 482
 - memory allocation, 633
 - pointers, invalidating, 491
 - pointers to
 - invalidating, 489–493
 - modifying, 487–489
 - storing, 468–473
 - setting, 631–633
 - size assumptions, 474–475
- envp environment pointer
- environ alternative, 491
- _environ alternative, 492–493
- invalidating, 491–493
- EOF, in-band errors, 537–540. *See also* End-of-file.
- Equal sign (=), equality operator, 283
- Equal signs (==), equality operator, 283
- ERANGE, 228, 558–561
- erfc(), 229
- errno error codes. *See also* Error handling; Return codes, checking.
 - alternative to, 535–537
 - checking, 558–563
 - file stream errors, checking for, 535–537
 - indeterminate values, 564–567
 - and library functions, 558–563
 - redefining, 563
 - return code type, 57–59
 - setting to zero, 558–563
 - uses for, 559–560
- errno.h, 563
- errno_t error codes, 57–59
- Error handling
 - 3Rs of survivability, 534
 - address arguments, 552, 555
 - adopting a policy, 533–535
 - altering control flow, 554–555
 - availability, measuring, 533
 - dependability, measuring, 533
 - detecting errors, 549–556
 - fault handling strategy, 534–535
 - fault tolerance, measuring, 533
 - file stream errors, 535–537
 - floating-point numbers, 218–224
 - functions
 - defined by TR 24731-1, 541–543
 - return codes, 118–119. *See also* Errno error codes; Errno_t error codes.
 - runtime-constraint handlers, 541–543
 - testing for error conditions, 57–59
 - global error indicators, 553–554, 555
 - in-band error indicators, 537–540
 - input/output. *See* Input/output, error handling.
 - longjmp(), 554–555
 - memory allocation, 355–357
 - mitigating effects, 534–535
 - recognition, 534
 - recovery, 534
 - reliability, measuring, 533
 - reporting errors, 549–556
 - required system qualities, 533
 - resistance, 534
 - return values, 551, 555. *See also* Errno error codes.
 - risk assessment summary, 532
 - rules and recommendations
 - related, 532
 - summary of, 531
 - runtime, 39
 - setjmp(), 554–555
 - survivability, 533–534
 - termination strategy. *See* Program termination.
 - vulnerabilities, avoiding, 534–535
- Error indicators. *See also* Errno error codes; Error handling.
 - end-of-file, 438–440
 - global, 553–554, 555
 - in-band
 - EOF, 537–540
 - exceptions, 539
 - null pointers, 538–539
- Errors of addition, 576–578
- Errors of omission, 14, 574–576
- Escape characters, in file names, 591
- Ettercap vulnerability, 423–424
- Evaluation order. *See* Order of operations.

- Exclamation, equal sign (!=), equality operator, 283
 - Exclusive file access, 456
 - Exclusive file locks, 456, 458
 - `exec1()`, 484
 - `exec1e()`, 484
 - `exec1p()`, 484
 - Execution character set, 590
 - Execution environment, 590
 - `execv()`, 484
 - `execve()`, 484
 - `execvp()`, 484
 - `execvp()`, 484
 - `exit()`, 544–545, 547–549
 - `_Exit()`, 546, 547
 - `EXIT_FAILURE`, 544–546
 - `EXIT_SUCCESS`, 544–545
 - `exp()`, 229
 - `exp2()`, 229
 - `expm1()`, 229
 - Expressions
 - `&&` (ampersands), logical AND, 96–98, 113
 - `||` (vertical bars), logical OR, 96–98, 113
 - array size, calculating, 95–96
 - arrays, incompatible types, 258–259
 - `assert()`, 122–123
 - assertions, side effects, 122–123
 - bit-field assumptions, 172–173
 - buffer overflows, pointer arithmetic, 107–109
 - casting away
 - `const`-qualification, 102–104
 - volatile qualifications, 123–124
 - constant values
 - assuming, 105–106
 - modifying, 102–104
 - `const`-qualification, casting away, 102–104
 - data types
 - bit-field alignment, 115–116
 - bit-field overlap, 116–117
 - incompatibilities, 114–117
 - integers vs. floating-point, 114–115
 - volatile qualifications, casting away, 123–124
 - with floating-point numbers, rearranging, 214–215
 - functions
 - call results, modifying, 129–131
 - error checking, 118–119
 - returned values, checking, 118–119
 - void return type, 118–119
 - magic numbers, 127
 - memory, uninitialized, 124–128
 - `NDEBUG` preprocessor symbol, 122
 - `offsetof()` macro, 135–137
 - order of operations
 - controlling with parentheses, 93–95
 - evaluation of subexpressions, 111–113
 - exceptions, 113
 - macros, 15
 - sequence points, 119–121
 - side effects, 111–113
 - pointer arithmetic, 107–109
 - pointers
 - converting, 131–133
 - null, dereferencing, 128–129
 - object alignment, 131–133
 - taking size of, 95–96
 - risk assessment summary, 92–93
 - rules and recommendations
 - related, 93
 - summary of, 91–92
 - sequence points
 - modifying function call results, 129–131
 - order of operations, 7, 119–121
 - side effects, 119–121
 - short circuit behavior, 96–98
 - side effects
 - in assertions, 122–123
 - sequence points, 119–121
 - `sizeof()` operator
 - operands, side effects, 104–105
 - type size, determining, 95–96, 109–111
 - variable size, determining, 109–111
 - struct member alignment, 98–100
 - structures
 - aligning in memory, 100–102
 - byte-by-byte comparison, 100–102
 - element offset, determining, 135–137
 - packed, 99
 - padding, 98–100
 - size, as sum of parts, 98–100
 - type size
 - determining, 95–96, 109–111
 - hard coding, 109–111
 - variables
 - size, determining, 109–111
 - uninitialized, referencing, 124–128
 - volatile qualifications, casting away, 123–124
- F**
- Failure Mode, Effects, Criticality Analysis (FMECA), xxiv
 - Fault handling strategy, 534–535
 - Fault tolerance, measuring, 533
 - `fclose()`, 439
 - `FD_CLOEXEC` flag, 452–453
 - `fdim()`, 229
 - `fopen()`, 385
 - `fenv.h`, 219
 - `feof()`, 437, 438–440
 - `ferror()`
 - alternative to `errno`, 535–537
 - file errors, detecting, 438–440
 - file stream errors, checking for, 535–537
 - `fflush()`
 - end-of-file indicator, 439
 - flushing file streams, 392, 444–446
 - `fgetc()`
 - end-of-file, detecting, 436, 438
 - return codes, checking, 388
 - `fgetpos()`
 - failure codes, 560
 - file positioning, 464–465
 - pushing back characters, 410–411
 - `fgets()`
 - alternative to `gets()`, 432–433
 - characters and strings,
 - unbounded copying, 308
 - failure, resetting strings, 446–448
 - new-line character, reading, 440–442
 - return codes, checking, 432–433
 - truncating null-terminated byte strings, 280–282

- fgetws()
 - failure, resetting strings, 446–448
 - new-line character, reading, 440–442
- Field width, 172
- File descriptors, 372–374, 450–454
- File extensions, header file names, 24
- File functions, temporary file names, 455
- File I/O functions, return code
 - checking, 386–389
- File links, 617–620
- File names. *See also* Renaming files.
 - (dash), as leading character, 591
 - accessing devices, 426–431
 - accessing files by, 372–374
 - carriage returns in, 591
 - case sensitivity, 591
 - creating, 457–458, 458–460, 461–462
 - disallowed characters, 591–593
 - escape characters, 591
 - hard coded, 457–458
 - naming conventions, 591–593
 - new-line characters, 591
 - patterns for, 459
 - punctuation characters, 591
 - selecting, 461
 - spaces, 591
 - special file access, 426–431
 - uniqueness
 - criteria for temporary files, 455
 - ensuring, 461–462
 - predictability, 455, 459
 - vulnerabilities, 457–461
- FILE objects, copying, 443–444
- FILE pointers, 372–374
- File positioning
 - alternating input/output, 444–446
 - error handling, 388–389
 - fseek() vs. rewind(), 398–399
 - valid values for, 464–465
- File scope, 70, 72–73, 87
- File streams
 - alternating input/output, 444–446
 - errors, checking for, 535–537
 - file positioning
 - alternating input/output, 444–446
 - error handling, 388–389
- fseek() vs. rewind(), 398–399
 - valid values for, 464–465
- flushing, 444–446
- stream arguments with side effects, 448–450
- text mode vs. binary mode, 411–413, 442–443
- Files
 - access control
 - advisory locks, 456
 - exclusive access, 456
 - exclusive locks, 456, 458
 - jails, 418–420
 - locking, 456
 - mandatory locks, 456
 - permissions, creating, 394–397
 - sandboxes, 418–420
 - shared locks, 456, 458
 - accessing
 - with file descriptors, 372–374
 - by file name, 372–374
 - with FILE pointers, 372–374
 - unintentionally, 374, 383–386, 463
 - closing, 450–454
 - creating
 - access permissions, 394–397
 - assumptions about, 383–387
 - hard-coded names, 457–459
 - mode strings, 407–408
 - in shared directories, 457–459
 - temporary, in shared directories, 454–463
 - errors, detecting, 438–440
 - file functions, temporary file names, 455
 - FILE pointers, 372–374
 - file positioning
 - alternating input/output, 444–446
 - error handling, 388–389
 - fseek() vs. rewind(), 398–399
 - valid values for, 464–465
 - header. *See* Header files.
 - identifying
 - by attributes, 389–394
 - by name, 372–374
 - including wrong one, 16–18
 - input/output. *See* Input/output, files.
 - open, removing, 399–401
 - opening
 - assumptions about, 383–387
 - mode parameter, 407–408
 - race conditions, 424–426
 - simultaneous multiple times, 424–426
 - orphaned, removing, 456–457
 - overwriting, unintentional, 383–386
 - renaming
 - portable behavior, 405–407
 - preserving destination file, 403–404, 406
 - race conditions, 404, 406
 - removing destination file, 404–405, 406
 - rename(), 403
 - rewinding, 398–399, 444–446
 - secure directories, 413–418
 - security vulnerabilities, 372–374
 - temporary
 - creating in shared directories, 454–463
 - criteria for, 455
 - naming, 455, 459
 - predictable names, 455, 459
 - removing, 456–457
 - vulnerabilities, 457–461
- TOCTOU race conditions
 - accessing devices as files, 429
 - checking file names, 378, 382
 - unintended file access, 374
- unique names
 - criteria for temporary files, 455
 - ensuring, 461–462
 - predictability, 455, 459
 - vulnerabilities, 457–461
- update mode, 444–446
- Flags
 - close-on-exec, 452
 - FD_CLOEXEC, 452–453
 - NDEBUG, 597
 - O, 83
 - O_CLOEXEC, 452–453
 - O_CREAT, 384, 458
 - O_EXCL, 384, 458
 - O_EXLOCK, 458
 - O_NOFOLLOW, 618
 - O_NONBLOCK, 427
 - O_SHLOCK, 458
 - packing, 172
 - SA_RESETHAND, 505, 507
 - struct member alignment, 98
- Flash vulnerability, 357
- Flexible array members, memory allocation, 358–360

- Floating literals, 45
- Floating-point exception, 219–220
- Floating-point numbers
 - binary floating-point, 215
 - bounds checking, math functions, 229
 - compiler variations, 213–214
 - conversions, 219
 - decimal floating-point, 215
 - demotion, 236–239
 - divide-by-zero errors, 219, 228
 - domain errors, 227–234
 - error handling, 218–224
 - expressions, rearranging, 214–215
 - fractions, 224
 - functions
 - calling with complex values, 226–227
 - math, 227–234
 - integer conversions, 234–239
 - vs. integers, 114–115
 - limitations, 212–214
 - as loop counters, 224–226
 - overflow, 228
 - precision, 213, 215–218
 - range errors, 227–234
 - risk assessment summary, 211–212
 - rounding, 215
 - rules and recommendations
 - related, 212
 - summary of, 211
 - SEH (structured exception handling), 222–223
- Flushing file streams, 444–446
- `fma()`, 229
- FMECA (Failure Mode, Effects, Criticality Analysis), xxiv
- `fmod()`, 229
- `fopen()`. *See also* Creating files.
 - creating files
 - access permissions, 395–396
 - assumptions about, 383–387
 - hard-coded names, 457–459
 - mode strings, 407–408
 - in shared directories, 457–459
 - temporary, 457–459
 - file identification, 372
 - opening files, 383–387
 - return codes, checking, 433–434, 561–562
- `fopen_s()`
 - creating files, 383–384, 395–396
 - file access permissions, 395–396
- `fork()`, 630–631
- Format strings
 - creating, 370–371
 - related vulnerabilities, 424
 - user input, 421–424
- Formatted I/O, 178–181
- Fortify Taxonomy: Software Security Errors*, 330–331
- `fpos_t`, 464–465
- `fprintf()`, 421–422, 513
- `fputc()`, 388, 436
- `fputs()`, 422, 439
- Fractions, floating-point numbers, 224
- `fread()`, 401–403
- `free()`
 - dangling pointers, 322–323
 - freeing nondynamic memory, 355–357
 - information leakage, 329–330
 - in signal handlers, 513
- Freed memory. *See* Memory, freed.
- Freeing memory. *See* Memory, freeing.
- `freopen()`, 372
- Friedl, Stephen, 44, 56, 602
- `fscanf()`, 157–158, 439
- `fseek()`
 - file positioning, 388–389, 398–399, 444–446
 - line breaks, 412
 - new-line character, 412
 - vs. `rewind()`, 398–399
- `fsetpos()`
 - failure codes, 560
 - file positioning, 444–446, 464–465
 - pushing back chars, 410–411
- `fstat()`, 390–394, 428, 634–635
- `ftell()`, 560
- Function declarators, type information, 51–54
- Function designator, 112–113
- Function prototypes, 52–53
- Function type, 51–54
- Function-like macros, 6–11, 328
- Functions
 - arguments
 - changing, 66
 - destination pointers, 80–82
 - passing by value, 66
 - referencing overlapping objects, 80–82
 - restrict qualification, 80–82
 - unintentional modification, 35
- asynchronous-signal-safe, 511–516
- calls
 - casting into pointers, 324–328
 - from macros, 30–32
 - results, modifying, 129–131
- characters and strings
 - alternatives, 288–291
 - arguments to, 314–315
- `errno` error codes, 57–59
- `errno_t` error code, 57–59
- error handling
 - defined by TR 24731-1, 541–543
 - return values, checking, 118–119
 - runtime-constraint handlers, 541–543
- floating-point numbers
 - calling with complex values, 226–227
 - math, 227–234
- inline, macro alternative, 6–11, 30–32
- inline substitution, 6
- input/output. *See* Input/output, functions.
- local, 10
- names, global replacement, 26–27
- parameters, 66–68
- pointers
 - referencing overlapping objects, 80–82
 - restrict qualification, 80–82
 - type conversion, 84–86
 - type information, declaring, 53–54
- program termination
 - `abort()`, 546–547, 547–549, 556–557, 564
 - `assert()`, 556–557
 - `atexit()`, 545
 - `exit()`, 544–545, 547–549
 - `_Exit()`, 546, 547
 - return from `main()`, 545–546, 547
 - summary of, 547
- replacing with less secure functions, 26–27
- returned values, checking, 118–119

- Functions, *continued*
 for signal handlers, 511–516
 static, macro alternative, 6–11, 30–32
 testing for error conditions, 57–59
 type-generic, 10–11
 variadic
 type alignment error, 63
 type interpretation error, 62–63
 type issues, 62–64
 va_arg() macro, 61, 62–64
 vfprintf(), 61
 vfscanf(), 61
 vscanf(), 61
 vsprintf(), 61
 vsscanf(), 61
 writer/caller contract, 59–62
 void return type, 118–119
 vsprintf(), noncompliant example, 26–27
- fwide(), 560
 fwrite(), 401–403
- G**
- getc()
 end-of-file, detecting, 436, 438
 return codes, checking, 388
 stream arguments with side effects, 448–450
- getchar()
 characters and strings,
 unbounded copying, 309
 end-of-file, detecting, 436, 438
 return codes, checking, 388
- getconf(), 481
- getenv()
 dangling pointers, 468
 environment variables
 clearing, 481–482
 copying, 298–299
 duplicate names, 475–478
 pointers from
 modifying, 487–489
 storing, 468–473
 race conditions, 468
 thread safety, 468
- getenv_s(), 469–470
 GetFileType(), 430
 GetFullPathName(), 381–382
- gets()
 alternative to, 432
 characters and strings, unbounded copying, 307–308
 deprecated use, 431
 truncating null-terminated byte strings, 280–282
- gets_s(), 309
- Global error indicators, 553–554, 555
- Greater-than signs (>>), right-shift operator, 174–175
- Guidelines. *See* Recommendations; Rules.
- H**
- Hacker's Delight*, 155
- Hand-coded code, 2–3
- Hard links, 617, 619
- Hard-coded file names, 457–458
- Header files. *See also* Files.
 inclusion guards, 21
 standard names
 (period), unique file names, 24
 case sensitivity, 24
 file extensions, 24
 list of, 17
 long file names, 25
 reusing, 16–18
 uniqueness, 17, 24–26
- Hiding
 information, 64–65
 objects, 70–72
- \$HOME value, 482
- HUGE_VAL, 228
- HUGE_VALF, 228
- HUGE_VALL, 228
- hypot(), 229
- I**
- ID numbers for vulnerabilities, xxvii
- Identifiers
 conflicting linkage classifications, 87–89
 duplicates, 78
 external linkage, 87–89
 implicit declaration, 4–77
 internal linkage, 87–89
 no linkage, 87–89
 significant characters, minimum requirements, 78
- uniqueness, 78–79
 visual distinction, 38–39
- ilogb(), 229
- Immutable objects, const-qualifying, 14, 35–36
- Immutable string literals, 284–285
- imod(), 169
- Implementation-defined nonportability, 603
- Implicit conversion, 324
- Implicit declaration, identifiers, 4–77
- In-band error indicators
 EOF, 537–540
 exceptions, 539
 null pointers, 538–539
- Including header files, 16–18, 21
- Incomplete array types, 243
- Incorrect assumptions, identifying, 39–42
- incr(), 350
- Increment side effects, 30–32
- Indices of arrays, range checking, 250–251
- Information hiding, 64–65
- Information leakage. *See also* Data loss.
 file descriptor leakage, 450–454
 from freed resources, 328–332
 improper permissions, 417
 multithreaded environments, 629
 uninitialized memory, 346–348
- Information Technology...Avoiding Vulnerabilities...*, xxv
- Initialization. *See also* Declarations and initialization.
 arrays, 243, 248–250
 character arrays, 312–313
 memory, 346–348
- Injection attacks, 279
- Inline functions, macro alternative, 6–11, 30–32
- Inline substitution, 6
- Input/output
 binary data, transferring across systems, 402–403
 character class ranges, 370–371
 conversion specifiers, 370–371
 data transfer across systems, 402–403
 denial-of-service attacks, 426
 devices, accessing as files, 426–431

- directory names, canonical form
 - conversion. *See* Canonical form conversion, path names.
- disk files, identifying, 430
- endianess, 401
- end-of-file
 - detecting, 436–438, 438–440
 - error indicators, 438–440
 - indicator, 436, 438
- fgets() failure, resetting strings, 446–448
- fgetws() failure, resetting strings, 446–448
- FILE objects, copying, 443–444
- format strings
 - creating, 370–371
 - related vulnerabilities, 424
- line breaks, 412, 440–442
 - input, 421–424
- new-line character
 - cross-environment differences, 412
 - reading, 440–442, 442–443
 - replacing, 442–443
- pushing back characters, 409–411
- reading data
 - binary data, 401–403
 - data type assumptions, 442–443
 - at link targets, 623–625
 - new-line character, 440–442, 442–443
 - restricting. *See* Access control.
- risk assessment summary, 368–369
- rules and recommendations
 - related, 370
 - summary of, 367–368
- writing data
 - binary, 401–403
 - to disk, 338–341
 - restricting. *See* Access control.
- Input/output, error handling
 - end-of-file errors, 438–440
 - file errors, detecting, 438–440
 - file positioning, 388–389
 - return codes, checking
 - character I/O functions, 436–438
 - file I/O functions, 386–389
 - I/O functions, 431–436
 - stream alteration, 408–409
- Input/output, file names. *See also* Renaming files.
 - accessing devices, 426–431
 - accessing files by, 372–374
 - canonical form conversion. *See* Canonical form conversion, path names.
 - checking. *See* Canonical form conversion, path names.
 - creating, 457–458, 458–460, 461–462
 - hard coded, 457–458
 - patterns for, 459
 - selecting, 461
 - special file access, 426–431
 - uniqueness
 - criteria for temporary files, 455
 - ensuring, 461–462
 - predictability, 455, 459
 - vulnerabilities, 457–461
- Input/output, file streams
 - alternating input/output, 444–446
 - file positioning
 - alternating input/output, 444–446
 - error handling, 388–389
 - fseek() vs. rewind(), 398–399
 - valid values for, 464–465
 - flushing, 444–446
 - stream arguments with side effects, 448–450
 - text mode vs. binary mode, 411–413, 442–443
- Input/output, files
 - access control
 - advisory locks, 456
 - exclusive access, 456
 - exclusive locks, 456, 458
 - jails, 418–420
 - locking, 456
 - mandatory locks, 456
 - permissions, creating, 394–397
 - sandboxes, 418–420
 - shared locks, 456, 458
 - accessing
 - with file descriptors, 372–374
 - by file name, 372–374
 - with FILE pointers, 372–374
 - unintentionally, 374, 383–386, 463. *See also* Access control.
 - closing, 450–454
 - creating
 - access permissions, 394–397
 - assumptions about, 383–387
 - temporary, in shared directories, 454–463
 - errors, detecting, 438–440
 - file descriptor leakage, 450–454
 - file descriptors, 372–374
 - file functions, temporary file names, 455
 - FILE pointers, 372–374
 - identifying
 - by attributes, 389–394
 - by name, 372–374
 - naming. *See* File names; Renaming files.
 - open, removing, 399–401
 - opening
 - assumptions about, 383–387
 - mode parameter, 407–408
 - race conditions, 424–426
 - simultaneous multiple times, 424–426
 - orphaned, removing, 456–457
 - overwriting, unintentional, 383–386
 - positioning
 - alternating input/output, 444–446
 - error handling, 388–389
 - fseek() vs. rewind(), 398–399
 - valid values for, 464–465
 - renaming. *See also* File names.
 - portable behavior, 405–407
 - preserving destination file, 403–404, 406
 - race conditions, 404, 406
 - removing destination file, 404–405, 406
 - rename(), 403
 - rewinding, 398–399
 - secure directories, 413–418
 - security vulnerabilities, 372–374
 - temporary
 - creating in shared directories, 454–463
 - criteria for, 455
 - naming, 455, 459
 - predictable names, 455, 459
 - removing, 456–457
 - vulnerabilities, 457–461

- Input/output, files, *continued*
 - TOCTOU race conditions
 - accessing devices as files, 429
 - checking file names, 378, 382
 - unintended file access, 374
 - unique names
 - criteria for temporary files, 455
 - ensuring, 461–462
 - predictability, 455, 459
 - vulnerabilities, 457–461
 - update mode, 444–446
- Input/output, functions
 - file functions, temporary file names, 455
 - file identification, 372–374
 - return codes, checking
 - character I/O functions, 436–438
 - file I/O functions, 386–389
 - I/O functions, 431–436
 - return values, 386–388
 - in signal handlers, 513
- Input/output, path names
 - canonical form conversion
 - Linux, 380–381
 - overview, 374–375
 - POSIX, 376–380
 - Windows, 381–382
 - simplifying. *See* Canonical form.
 - validating, 374–383
- Input/output, race conditions
 - opening files, 424–426
 - renaming files, 404, 406
 - shared files, 456
- TOCTOU
 - accessing devices as files, 429
 - checking file names, 378, 382
 - unintended file access, 374
 - unlinking files, 463
- int type, 274, 436–438
- Integer constant expression, 46
- Integer conversion rank, 150
- Integer literals, 45, 48–49
- Integer overflow
 - memory allocation, 362–363
 - signed integers
 - addition, 192–194
 - division, 197
 - left-shift operations, 199–200
 - modulo operation, 197–198
 - multiplication, 195–196
 - overview, 191–192
 - subtraction, 194–195
 - unary negation, 198–199
 - vulnerabilities, 146
- Integer promotions, 149–150, 152
- Integer type conversions. *See also* Casting; Conversions.
 - data loss, 186–191
 - integer-to-pointer, 170–172
 - loss of precision, 188–189
 - minimum ranges, 189–190
 - misinterpreted data, 186–191
 - numeric strings, to greatest-width integers, 143
 - pointer-to-integer, 170–172
 - rules for
 - arithmetic conversions, 150–151
 - code examples, 151–153
 - integer conversion rank, 150
 - integer promotions, 149–150, 152
 - overview, 149
 - required conversions, determining, 150
 - usual arithmetic conversions, 150–151
 - signed to signed, 188–189
 - signed to unsigned, 188
 - string data
 - with functions, 157–162
 - to greatest-width integers, 143
 - unsigned to signed, 187
 - unsigned to unsigned, 188–189
- Integers
 - % (percent sign), remainder operator, 168–170
 - addition
 - integer overflow, 192–194
 - unsigned integer wrapping, 182–183
 - assigning to a larger size, 207–210
 - bit-field assumptions, 172–173
 - bitwise operations
 - combined with arithmetic operations, 175–178
 - on unsigned operands, 174–175
 - char type, 162–164
 - character types, 162–164
 - comparing to a larger size, 207–210
 - data models
 - code examples, 143–145
 - definition, 141
 - integral ranges, 142
 - list of, 142
 - numeric strings, converting to
 - greatest-width integers, 143
 - size restrictions, 142–143
 - types available, 142–143
 - division
 - divide-by-zero errors, 201
 - fractional remainders. *See* Modulo operations.
 - integer overflow, 197
 - by powers of 2, 176–177
 - remainder assumptions, 168–170
 - with shift operators, 176–177
 - truncation toward zero, 169
 - enumeration constants, mapping, 167–168
 - vs. floating-point, 114–115
 - floating-point conversions, 234–239
 - formatted I/O, 178–181
 - IntergLib, 154–155
 - left shift
 - multiplication with, 176–177
 - negative numbers, 203–207
 - unsigned integer wrapping, 185
 - modulo arithmetic, 165
 - modulo operations
 - divide-by-zero errors, 201–203
 - division, 169
 - integer overflow, 197–198
 - modwrap semantics, 164–165
 - multiplication
 - integer overflow, 195–196
 - by powers of 2, 176–177
 - with shift operators, 176–177
 - unsigned integer wrapping, 184
 - numeric values, character types
 - for, 162–164
 - object size, representing, 145–149
 - range checking, 153–157, 164–166
 - remainder assumptions, 168–170. *See also* Modulo operations.
 - restricted range usage, 165
 - right shift
 - division with, 177
 - negative numbers, 203–207
 - on unsigned operands, 174–175
 - risk assessment summary, 140–141

- rsize_t type, 145–149
 - rules and recommendations
 - related, 141
 - summary of, 139–140
 - saturation semantics, 164–165
 - secure integer libraries, 153–155
 - signed char type, 162–164
 - signed integer overflow, 146
 - size assumptions, 141
 - size_t type, 145–149
 - subtraction
 - integer overflow, 194–195
 - unsigned integer wrapping, 183–184
 - type definition, 178–181
 - unsigned char type, 162–164
 - unsigned integer wrapping, 181–186
 - from untrusted sources, 155–157
 - wrapping, 164–165, 181–186
 - Integer-to-pointer conversions, 170–172
 - IntergLib, 154–155
 - Internal linkage identifiers, 87–89
 - Internet resources. *See* Online resources.
 - Interruptible signal handlers, signal() calls, 526–529
 - intmax_t, 178–181
 - inttypes.h, 143
 - I/O functions. *See* Input/output, functions.
 - isalnum(), 314
 - isalpha(), 314
 - isascii(), 314
 - isblank(), 314
 - isctrnl(), 314
 - isgraph(), 314
 - islower(), 314
 - isprint(), 314
 - ispunct(), 314
 - isspace(), 314
 - isupper(), 314
 - isxdigit(), 314
 - Iterating through arrays, 259–261. *See also* Loops.
- J**
- Jails, 418–420
- K**
- Kerberos 5 vulnerability, 319–321
 - Kettlewell, Richard, 431
- L**
- Lawrence Livermore National Laboratory, xx
 - ldexp(), 229
 - Leaking information. *See* Data loss, information leakage.
 - Leffler, Jonathan, 18, 86, 346, 482
 - Left shift
 - multiplication with, 176–177
 - negative numbers, 203–207
 - unsigned integer wrapping, 185
 - Length modifier, 370
 - lessen_memory_usage(), 302–303
 - Less-than signs (<<), left-shift operator, 174–175
 - lgamma(), 229
 - Libraries, null-terminated byte strings, 275–276
 - Library functions, 558–563, 588–590
 - Lifetime, objects, 72–74
 - limits.h, 142
 - Line breaks, 412, 440–442
 - Linkage classifications, conflicting, 87–89
 - Links
 - checking for, 617–620
 - file, 617–620
 - hard, 617, 619
 - reading target of, 623–625
 - soft, 617
 - symbolic, 617
 - virtual drives, 617
 - Linux kernel vmsplICE exploit, 186
 - Literals
 - declaring and initializing. *See* Declarations and initialization, literals.
 - string. *See* Characters and strings, string literals.
 - Local functions, 10
 - lockfile(), 456
 - lockfileEx(), 456
 - Locking
 - files
 - advisory locks, 456
 - exclusive access, 456
 - exclusive locks, 456, 458
 - jails, 418–420
 - mandatory locks, 456
 - shared locks, 456, 458
 - memory, 339
 - log(), 229
 - log1p(), 229
 - log2(), 229
 - log10(), 229
 - logb(), 229
 - Logical completeness, 572–574
 - Long double, 237
 - Long file names, header files, 25
 - long int
 - integer conversion rank, 150
 - smallest possible value, 142
 - long long int, 150
 - longjmp()
 - atexit() handler calls, terminating, 494
 - calls from signal handlers, 519–523
 - error handling, 554–555
 - invalid calls to, 496–497
 - in signal handlers, 519–523
 - Loops
 - \ (backslash), loop terminator, 260
 - buffer overflows, 259–261
 - counters, with floating-point numbers, 224–226
 - do-while, 27–29
 - freeing memory, 353
 - looping beyond last element, 259–261
 - optimizing away, 83
 - terminating, 82, 259–261
 - Loss of data. *See* Data loss.
 - lrint(), 229
 - lround(), 229
 - lstat(), 428, 634
 - lstat-fopen-fstat idiom, 618–619
- M**
- Macro definitions vs. type definitions, 15–16
 - Macro parameters
 - # (pound sign) operator, concatenating tokens, 12, 18–20
 - ## (pound signs) operator, stringification, 12, 18–20
 - concatenating string literals, 12
 - concatenating tokens, 12
 - names, parenthesizing, 11–13
 - names, surrounded by commas, 12
 - stringification, 12, 18–20

- Macro replacement
 - # (pound sign) operator, concatenating tokens, 12, 18–20
 - ## (pound signs) operator, stringification, 12, 18–20
 - concatenating tokens, 12, 18–20
 - stringification, 12, 18–20
 - token pasting, 19
- Macros
 - alternatives to, 6–11, 30–32
 - assignment side effects, 30–32
 - call-by-name semantics, 11
 - compile-time constant, yielding, 10
 - decrement side effects, 30–32
 - function calls, 30–32
 - function names, global replacement, 26–27
 - function-like, 6–11, 328
 - increment side effects, 30–32
 - inline functions, 6–11, 30–32
 - local functions, 10
 - multi-statement, in `do-while` loops, 27–29
 - operator precedence, 15
 - operator precedence in expressions, 15
 - replacement lists, parenthesizing, 13–15
 - side effects, avoiding, 30–32
 - static functions, 6–11, 30–32
 - type-generic functions, 10–11
 - unsafe, 30–32
 - volatile access side effects, 30–32
- Magic numbers
 - buffer overflow, 50
 - in expressions, 127
 - readability, 45
- `main()`, 545–546
- `malloc()`
 - environment variables, copying, 471–473
 - memory allocation errors, 355–357
 - memory size, calculating, 362–365
 - returned pointers, casting, 324–328
 - zero-length allocations, 333
- Managed strings, 291–293
- Mandatory file locks, 456
- Masking signals, 500–503
- Math functions, 227–234
 - `MATH_ERREXCEPT`, 228
 - `MATH_ERRNO`, 228
 - `math.h`, 226, 227, 559
 - `mbstate_t`, 464
 - `mbstowcs()`
 - comparing return values, 611
 - copying overlapping objects, 81
 - Meaningful literal names, 45–51
 - Member alignment, 98–100
 - `memcmp()`, 101
 - `memcpy()`
 - copying overlapping objects, 81
 - memory allocation size, range checking, 363–364
 - null-terminated byte strings, memory allocation, 297
 - `memmove()`, 81
- Memory
 - allocating
 - buffer overflows, 343
 - environment variables, 633
 - error handling, 355–357
 - for flexible array members, 358–360
 - function calls, casting into pointers, 10, 324–328
 - integer overflow, 362–363
 - invalid size arguments, 365
 - out-of-bounds check, 342–343
 - range checking, 363–364
 - at same abstraction level, 319–321
 - in same modules, 319–321
 - size, calculating, 362–365
 - stack allocations, 335–338
 - wrapping, checking for, 342–343
 - zero-length space, 332–335
 - buffer overflows
 - invalid size arguments, 365
 - memory initialization, 346–348
 - unsigned integer wrapping, 343
 - core dumps, 338–339
 - dangling pointers, 322–323, 351–353
 - denial-of-service attacks
 - allocation failures, 355–357
 - memory management errors, 321
 - reading freed memory, 352
 - stack allocation, 337–338
 - freed
 - accessing, 351–353
 - writing to, 321, 323
 - freeing
 - access-freed-memory vulnerability, 322–323
 - clearing returned resources, 328–332, 582–585
 - double-free vulnerability, 319–321, 322–323, 353–355
 - nondynamic memory, 360–362
 - at same abstraction level, 319–321
 - in same modules, 319–321
 - in signal handlers, 513
 - information leakage, 328–332, 346–348
 - initializing, 346–348
 - literal requirements, 46–47
 - locking, 339
 - overwrites, array errors, 251, 253
 - paging, 338–340
 - pointer conversions, unintended, 328
 - pointer validation, 348–351
 - recursion, and stack allocation, 336–337
 - risk assessment summary, 318
 - rules and recommendations related, 318–319
 - summary of, 317–318
 - sensitive information
 - clearing, 328–332
 - leaking, 328–332, 346–348
 - writing to disk, 338–341
 - swapping. *See* Paging.
 - uninitialized, referencing, 124–128
 - usage, reducing, 301–303
- `memset()`
 - casting away `const` qualification, 103
 - compiler optimization, 330–332, 582–585
 - string truncation, 300
 - unintended write to memory, 108–109
- `memset_s()`, 584
- Messier, Matt, 277, 397
- Metrics
 - risk assessment, xxiv
 - vulnerabilities, xxvi
- Mitigating error effects, 534–535
- MITRE, xxiv
- `mkstemp()`, 455, 461

- mktemp()
 - creating temporary files, 459–460
 - deprecated use, 460
 - temporary file criteria, 455
- mlock(), 339–340, 341
- mode parameter, 407–408
- Modernizing Legacy Systems*, xxxi
- Modules, memory allocation, 319–321
- Modulo arithmetic, 165
- Modulo operations
 - divide-by-zero errors, 201–203
 - division, 169
 - integer overflow, 197–198
- Modwrap semantics, 164–165
- Morris worm, xv
- Mozilla SVG heap buffer wrap, 184
- Multibyte character strings, 273
- Multiplication
 - integer overflow, 195–196
 - by powers of 2, 176–177
 - with shift operators, 176–177
 - unsigned integer wrapping, 184
- Multithreaded environments, 626–629
- Mutable string literals, 284–285
- Mutexes
 - bit-fields in multithreaded environments, 626–629
 - destroying, 625–626
 - ownership, 625–626
 - race conditions, 627
 - unlocking, 625–626
- N**
- Named symbolic literals, 45–47, 50
- Names
 - of files. *See* File names; Renaming files.
 - of vulnerabilities, xxvii
- NaN, 224
- Narrow character strings, 273, 284
- National use positions, 590
- NDEBUG, 122, 597
- Negation operator, 198
- New-line character
 - cross-environment differences, 412
 - in file names, 591
 - reading, 440–442, 442–443
 - replacing, 442–443
- nextafter(), 229
- nexttoward(), 229
- No linkage identifiers, 87–89
- Noncompliant code examples, xxiii
- Nondynamic memory, freeing, 360–362
- Noninterruptible signal handlers, masking signals, 500–503
- Nonpersistent signal handlers, 506–507, 526
- Non-prototype-format declarators, 51–52
- Notation consistency, arrays, 251–254
- Null pointers, 128–129, 538–539
- Null termination character, 299–303
- Null wide character, 273
- Null-terminated byte strings
 - buffer overflows, 280–282, 294–299
 - creating, 299–303
 - definition, 273
 - dynamic allocation, 275–276
 - environment variables, copying, 298–299
 - libraries, 275–276
 - managing consistently, 275–276
 - null termination character, 299–303
 - off-by-one error, 294–295
 - static allocation, 275–276
 - storage allocation, 294–299
 - traversing arrays, 259–261
 - truncating, 280–282, 300–301
 - unintended character arrays, 312–313
- Numeric conversion, 559
- Numeric literals, buffer overflow, 50
- Numeric strings, converting to greatest-width integers, 143
- Numeric values, character types for, 162–164
- O**
- 0 flag, 83
- Object-like macros, 47
- Objects
 - access restrictions, 82–84
 - accessing outside their lifetime, 72–74
 - alignment, 131–133
 - caching restrictions, 82–84
 - confined to current scope, 70–72
 - declaring as `static`, 70–72
 - hiding, 70–72
 - lifetime duration, 72–74
 - overlapping, referencing, 80–82
 - size, representing, 145–149
 - volatile-qualified type, 82–84
- `O_CLOEXEC` flag, 452–453
- `O_CREAT` flag, 384, 458
- `O_EXCL` flag, 384, 458
- `O_EXLOCK` flag, 458
- Off-by-one error, 294–295
- `offsetof()` macro, 135–137
- Online resources
 - CERT Secure Coding Standards wiki, xvii
 - code analysis tools, xx
 - IntergLib, 154–155
 - Lawrence Livermore National Laboratory, xx
 - secure integer libraries, 154–155
- `O_NOFOLLOW` flag, 618
- `O_NONBLOCK` flag, 427
- Opaque data types, 64–65
- `open()`
 - creating files
 - access permissions, 396–397
 - hard-coded names, 457–459
 - in shared directories, 457–459
 - temporary, 457–459, 459–460
 - file permissions, 396–397
- Open files, removing, 399–401
- OpenBSD operating system flaw, 108
- Opening files
 - assumptions about, 383–387
 - mode parameter, 407–408
 - race conditions, 424–426
 - simultaneous multiple times, 424–426
- OpenSSL package, Debian Linux, 128
- Optimization, and sensitive information, 582–585
- OR operators
 - `^` (caret), bitwise OR, 174–175
 - `|` (vertical bar), bitwise inclusive OR, 174–175
 - `||` (vertical bars), logical OR, 96–98, 113
- Order of operations in expressions
 - `&&` (ampersands), logical AND, 113
 - `||` (vertical bars), logical OR, 113
 - controlling with parentheses, 93–95
 - evaluation of subexpressions, 111–113

- Order of operations in expressions, *continued*
 - exceptions, 113
 - macros, 15
 - sequence points, 119–121
 - side effects, 111–113
- Orphaned files, removing, 456–457
- O_SHLOCK flag, 458
- Out-of-bounds check, memory allocation, 342–343
- Overflow
 - buffers. *See* Buffer overflow.
 - floating-point numbers, 228
 - integers. *See* Integer overflow.
- Overwriting data, 473
- Overwriting files, 383–386
- Ownership, mutexes, 625–626

- P**
- Packed structures, 99
- Padding structures, 98–100
- Paging, 338–340
- Parenthesizing
 - macro names, 11–13
 - macro replacement lists, 13–15
- Path names
 - canonical form conversion
 - Linux, 380–381
 - overview, 374–375
 - POSIX, 376–380
 - Windows, 381–382
 - validating, 374–383
- PATH_MAX
 - file system variance, 379
 - output buffer size, 376–379
 - value, obtaining, 376, 379–380
- Patterns for file names, 459
- Percent sign (%), remainder operator, 168–170
- Period (.), unique header file names, 24
- Permission escalation
 - signal handlers, 516
 - temporary file creation, 463
 - TOCTOU race conditions, 635
- Permissions
 - appropriate level of, 394–397
 - improper, information leakage, 417
 - principle of least privilege, 620–623
- perror(), 560
- Persistent signal handlers, 503–507, 526–529
- Plain char type, 274, 282–284
- Platform dependencies, avoiding, 602–604
- Pointer arithmetic, 107–109
- Pointers
 - arrays
 - adding, 263–269, 265–269
 - arithmetic operations on, 261–263, 263–265, 265–269
 - comparing, 261–263
 - subtracting, 261–263, 263–265, 265–269
 - to characters and strings, 273
 - conversions, unintended, 328
 - dangling
 - accessing freed memory, 351–353
 - definition, 351
 - free(), 322–323
 - getenv(), 468
 - to environment variables
 - invalidating, 489–493
 - modifying, 487–489
 - storing, 468–473
 - in expressions
 - converting, 131–133
 - null, dereferencing, 128–129
 - object alignment, 131–133
 - taking size of, 95–96
 - memory allocation calls, casting into, 10
 - null, in-band error indicators, 538–539
 - validation, 348–351
 - to values, function parameters, 66–68
- Pointer-to-integer conversions, 170–172
- popen(), 482
- Portability
 - implementation-defined nonportability, 603
 - platform dependencies, avoiding, 602–604
 - special characters, 590
 - unspecified nonportability, 603
 - vfork(), 629
- Positioning files. *See* File positioning.

- POSIX
 - access permissions
 - principle of least privilege, 620–623
 - relinquishing, 636–642
 - buffer overflow, readlink()
 - termination, 624
 - denial-of-service attacks, vfork(), 629
 - environment variables, 631–633
 - links, 617–620, 623–625
 - mutexes
 - bit-fields in multithreaded environments, 626–629
 - destroying, 625–626
 - ownership, 625–626
 - race conditions, 627
 - unlocking, 625–626
 - portability issues, 629
 - principle of least privilege, 620–623
 - putenv(), 631–633
 - race conditions, 615–616, 633–636
 - readlink(), 623–625
 - risk assessment summary, 614
 - rules and recommendations
 - related, 614–615
 - summary of, 613–614
 - security issues, vfork(), 629
 - setenv(), 633
 - vfork(), 629–631
- Pound sign (#) operator, concatenating tokens, 12, 18–20
- Pound signs (##) operator, stringifying macro parameters, 12, 18–20
- pow(), 229, 233–234
- Precision
 - floating-point numbers, 213, 215–218
 - loss of, integer conversions, 188–189
- Predictability, file names, 455, 459
- Preprocessors
 - ? (question mark), repeating, 22–24
 - concatenation
 - # (pound sign) operator, concatenating tokens, 12, 18–20
 - string literals, 12
 - token pasting, 19

- tokens, 12, 18–20, 29–30
 - universal character names, 29–30
 - files. *See also* Header files.encoding types, macro definitions vs. type definitions, 15–16 including wrong one, 16–18 inclusion guards, 21
 - functions
 - calls, from macros, 30–32
 - inline, macro alternative, 6–11, 30–32
 - local, 10
 - names, global replacement, 26–27
 - replacing with less secure functions, 26–27
 - static, macro alternative, 6–11, 30–32
 - type-generic, 10–11
 - immutable objects, `const-qualifying`, 14
 - inline functions, macro alternative, 6–11
 - macro definitions vs. type definitions, 15–16
 - macro parameters
 - `#` (pound sign) operator, concatenating tokens, 12, 18–20
 - `##` (pound signs) operator, stringification, 12, 18–20
 - concatenating string literals, 12
 - concatenating tokens, 12
 - names, parenthesizing, 11–13
 - names, surrounded by commas, 12
 - stringification, 12, 18–20
 - macro replacement
 - `#` (pound sign) operator, concatenating tokens, 12, 18–20
 - `##` (pound signs) operator, stringification, 12, 18–20
 - concatenating tokens, 12, 18–20
 - stringification, 12, 18–20
 - token pasting, 19
 - macros
 - alternatives to, 6–11, 30–32
 - call-by-name semantics, 11
 - compile-time constant, yielding, 10
 - function calls, 30–32
 - function names, global replacement, 26–27
 - inline functions, 6–11, 30–32
 - local functions, 10
 - multi-statement, in `do-while` loops, 27–29
 - operator precedence, 15
 - replacement lists, parenthesizing, 13–15
 - side effects, avoiding, 30–32
 - static functions, 6–11, 30–32
 - type-generic functions, 10–11
 - unsafe, definition, 30
 - unsafe, invoking, 30–32
 - memory allocation function calls, casting into pointers, 10
 - risk assessment summary, 6
 - rules and recommendations
 - related, 6
 - summary of, 5
 - static functions, macro alternative, 6–11
 - three-character sequences, 22–24
 - trigraph sequences, 22–24
 - type definitions
 - vs. macro definitions, 15–16
 - scope rules, 15
 - `typedef` vs. `#define`, 15–16
 - universal character names, creating, 29–30
 - variables, multiple per declaration, 16
 - Preprocessors, header files. *See also* Files.
 - inclusion guards, 21
 - standard names
 - `.` (period), unique file names, 24
 - case sensitivity, 24
 - file extensions, 24
 - list of, 17
 - long file names, 25
 - reusing, 16–18
 - uniqueness, 17, 24–26
 - Principle of least privilege, 620–623
 - `printf()`, 179
 - Privileges. *See* Access control.
 - Program termination
 - `abort()`, 546–547, 547–549, 556–557, 564
 - `assert()`, 556–557
 - `atexit()`, 494–497, 545
 - `exit()`, 544–545, 547–549
 - `_Exit()`, 546, 547
 - functions, summary of, 547
 - return from `main()`, 545–546, 547
 - sensitive information vulnerability, 549, 557
 - Pseudorandom number generation, 607–610
 - `ptrdiff_t`, 179
 - Punctuation characters, in file names, 591
 - Pushing back characters, 409–411
 - `putc()`
 - end-of-file, detecting, 436
 - return codes, checking, 388
 - stream arguments with side effects, 449–450
 - `putchar()`, 388, 436
 - `putenv()`, 631–633
 - `_putenv_s()`, 491–492
 - `puts()`, 439
- Q**
- Question mark (?), repeating, 22–24
- R**
- Race conditions
 - multiple threads, 615–616
 - mutexes, 627
 - opening files, 424–426
 - permission escalation, 635
 - renaming files, 404, 406
 - shared files, 456
 - signal handlers, 517–519, 528
 - symbolic links, 633–636
 - TOCTOU
 - accessing devices as files, 429
 - checking file names, 378, 382
 - unintended file access, 374
 - unlinking files, 463
 - `raise()`, 523–526
 - `rand()`, 607–610
 - `random()`, 608–609
 - Random number generation, 607–610
 - Range. *See also* Bounds.
 - checking
 - array indices, 250–251
 - array size arguments, 254–255
 - floating-point numbers, 227–234

- Range, *continued*
 checking, *continued*
 integers, 153–157, 164–166
 memory allocation, 363–364
 restricting, integers, 165
- Readability of code
 combining bitwise and arithmetic operations, 175
 comments, 42–43, 579–580
 literals, 45
 magic numbers, 45
 meaningful literal names, 45–51
 type definitions, 44–45
 typedef, 44–45
 variables, one per declaration, 42–44
- Reading
 data
 binary data, 401–403
 data type assumptions, 442–443
 at link targets, 623–625
 new-line character, 440–442, 442–443
 restricting. *See* Access control.
 freed memory, 352
- readlink(), 623–625
- realloc()
 clearing freed space, 331
 freeing nondynamic memory, 355–357
 information leakage, 330–332
 memory allocation errors, 355–357
 memory size, calculating, 362–365
 reducing memory usage, 301–303
 resizing dynamically allocated arrays, 343–346
 returned pointers, casting, 324–328
 zero-length allocations, 334–335
- realpath(), 376–380
- Recognition, 534
- Recommendations. *See also* Rules; *specific topics*.
 categories, xix. *See also specific categories*.
 definition, xix
 deviation procedures, 4
 identifiers, xxiii
 priorities and levels, xxiv, xxv, 4
 risk assessment priority, xxiii
 secure coding compliance, 4
 tool-generated, tool-maintained code, 3
- Recovery, 534
- Recursion, and stack allocation, 336–337
- References
Building Systems from Commercial Components, xxxi
 CERT/CC VND (CERT/CC Vulnerability Notes Database), xxiv
 CVE (Common Vulnerabilities and Exposures), xxiv
 CWE (Common Weakness Enumeration) IDs, xxiv
Delivering Signals for Fun and Profit, 519
 description, xxiv–xxv
Fortify Taxonomy: Software Security Errors, 330–331
 getting information about, xxiv–xxv
Hacker's Delight, 155
Information Technology...Avoiding Vulnerabilities..., xxv
Modernizing Legacy Systems, xxxi
Secure Coding Guide, 431
Secure Coding in C and C++, xxxi
Source Code Analysis Tool Functional Specification, 331
- Related Vulnerabilities sections, xxv–xxvii
- Relationships between literals, encoding, 54–56
- Reliability, measuring, 533
- Relinquishing access permissions, 636–642
- Remainder assumptions, 168–170
- Remainder operator. *See* Modulo operations.
- Remote Procedure Call (RPC), 259–260
- remove(), 372–373, 399–401
- Removing files
 open files, 399–401
 orphaned files, 456–457
 renamed destination files, 404–405, 406
 temporary files, 456–457
 unlinking, 486
- rename()
 file identification, 372
 portability, 405–407
 preserving destination files, 403–404, 406
 prototype, 403
 removing destination files, 404–405, 406
- Renaming files. *See also* File names.
 portable behavior, 405–407
 preserving destination file, 403–404, 406
 race conditions, 404, 406
 removing destination file, 404–405, 406
- Replacement lists, parenthesizing, 13–15
- Reporting errors, 549–556
- Resistance, 534
- restrict qualification, 80–82
- Return codes, checking. *See also* Errno error codes.
 character I/O functions, 436–438
 fgetc(), 388
 fgets(), 432–433
 file I/O functions, 386–389
 fopen(), 433–434
 fputc(), 388
 functions, 118–119
 getc(), 388
 getchar(), 388
 I/O functions, 431–436
 putchar(), 388
 putchar(), 388
 snprintf(), 432, 434–437
 stream alteration, 408–409
 ungetc(), 388
- Return from main(), 545–546, 547
- Return values
 comparing, 610–613
 data type, 610–613
 error handling, 551, 555
- Reusing header file names, 16–18
- rewind(), 398–399, 444–446
- Rewinding files, 398–399, 444–446
- Right shift
 division with, 177
 negative numbers, 203–207
 on unsigned operands, 174–175
- Risk assessment
 description, xxiii
 metrics, xxiv
 priorities and levels, xxiv, xxv

- summaries
 - arrays, 242
 - characters and strings, 272
 - declarations and initialization, 34–35
 - environment, 467–468
 - error handling, 532
 - expressions, 92–93
 - floating-point numbers, 211–212
 - input/output, 368–369
 - integers, 140–141
 - memory, 318
 - POSIX, 614
 - preprocessors, 6
 - signals, 499
- Rounding floating-point numbers, 215
- RPC (Remote Procedure Call), 259–260
- `rsize_t` type
 - array indices, range checking, 250–251
 - for integer values, 145–149
 - memory allocation size, calculating, 363–364
- Rules. *See also* Recommendations; *specific topics*.
 - categories, xviii. *See also specific categories*.
 - definition, xviii
 - deviation procedures, 4
 - identifiers, xxiii
 - priorities and levels, xxiv, xxv, 4
 - risk assessment priority, xxiii
 - secure coding compliance, 4
 - tool-generated, tool-maintained code, 3
- Runtime-constraint handlers, 541–543
- S**
- Saks, Dan, 16
- Samba vulnerability, 423–424
- Sanitizing
 - ASCII characters, 590–593
 - characters and strings, 276–280, 292, 590–593
 - data, 276–280, 292
 - the environment, 478–482
 - special characters, 276–280
- `SA_RESETHAND` flag, 505, 507
- Saturation semantics, 164–165
- `scalbln()`, 229
- `scalbn()`, 229
- `scanf()`
 - converting character strings to integers, 157–158
 - end-of-file indicator, 440
 - programmer-defined integer types, 180
- Scope
 - block, 42, 72–73
 - conflicting linkage classifications, 87–89
 - const-qualified objects, 45–46
 - file
 - assertions, 42
 - external linkage, 70–71
 - object storage duration, 72–73
 - specifier, 70–71
 - identifier declaration, 87–89
 - rules for type definitions, 15
- Secure Coding Guide*, 431
- Secure Coding in C and C++*, xxxi
- Secure development process, xvi–xvii
- Secure directories, 413–418
- Secure integer libraries, 153–155
- Secure products, lack of demand for, xvi
- `secure_dir()`, 414–417
- `SecureZeroMemory()`, 583
- SEH (structured exception handling), 222–223
- Sensitive information loss. *See also*
 - Data loss.
 - clearing, 328–332
 - compiler optimization, 582–585
 - improper permission vulnerability, 417
 - leaking. *See* Information leakage.
 - program termination vulnerability, 549, 557
 - writing to disk, 338–341
- Sequence points
 - modifying function call results, 129–131
 - order of evaluation, example, 7
 - order of operations, 119–121
 - side effects, 119–121
- `setbuf()`, 408–409
- `set_constraint_handler_s()`, 541–543
- `setenv()`
 - environment variables
 - duplicate names, 475–478
 - modifying, 488–489
 - pointers, invalidating, 490–491
 - heap memory, allocating, 633
- `setfile()`, 588–590
- `_set_invalid_parameter_handler()`, 543
- `setjmp()`, 554–555
- `setlocale()`, 561
- `setrlimit()`, 338–339
- `setvbuf()`, 408–409
- Shared file locks, 456, 458
- Shift operations
 - left shift
 - multiplication with, 176–177
 - negative numbers, 203–207
 - unsigned integer wrapping, 185
 - right shift
 - division with, 177
 - negative numbers, 203–207
 - on unsigned operands, 174–175
- Short circuit behavior, 96–98
- Side effects
 - assertions, 122–123
 - macros, 30–32
 - order of operations, 111–113
 - sequence points, 119–121
 - `sizeof()` operator, 104–105
- `sigaction()`
 - in interruptible signal handlers, 527–528
 - masking signals, 502–503
 - `raise()` calls from signal handlers, 524–526
 - signal handler interruptions, 502–503
 - signal handler persistence, 505–507
- `SIG_ERR`, 560, 564
- `signal()`
 - deprecated use, 506, 526, 528
 - `errno`, indeterminate values, 564–567
 - in interruptible signal handlers, 526–529
 - masking signals, 500–503
 - registering signal handlers, 500–503
 - signal handler persistence, 503–507
- Windows vs. UNIX, 503–504

- Signal handlers
 - asynchronous-signal-safe functions, 511–516. *See also* `longjmp()`, in signal handlers.
 - freeing memory, 513
 - functions, calling, 511–516
 - input/output, functions, 513
 - interruptible, `signal()` calls, 526–529
 - `longjmp()` calls, 519–523
 - noninterruptible, masking signals, 500–503
 - nonpersistent, 506–507, 526
 - persistent, 503–507, 526–529
 - race conditions, 517–519
 - recursive `raise()` calls, 523–526
 - registering, 500–503
 - shared objects, accessing or modifying, 517–519, 564–567
 - vulnerabilities, 516, 522
 - Windows vs. UNIX, 503–504
- Signals
 - implementing normal functionality, 507–511
 - masking, 500–503
 - noninterruptible signal handlers, 500–503
 - race conditions, 528
 - risk assessment summary, 499
 - rules and recommendations related, 500
 - summary of, 499
- signed char type, 162–164, 274, 305–307
- Signed integer conversions, 188–189
- Significant characters, minimum requirements, 78
- Simplifying path names. *See* Canonical form conversion, path names.
- Single-byte character strings, 273
- `sinh()`, 229
- Size
 - arrays
 - arguments, variable length arrays, 254–255
 - bounds specification, 312–313
 - determining, 245–249
 - insufficient, 255–258
 - string literal initialization, 312–313
 - environment variables, assumptions, 474–475
 - integers, assumptions, 141
 - memory allocation, calculating, 362–365
- `sizeof()`
 - array size, determining, 245–247
 - end-of-file, detecting, 438–440
 - operands, side effects, 104–105
 - signed integer overflow, 146
 - type size, determining, 49, 95–96, 109–111
 - variable size, determining, 49, 109–111
- `size_t` type
 - array indices, range checking, 250–251
 - `calloc()` arguments, 342–343
 - comparing return values, 611
 - for integer values, 145–149
 - memory allocation size, calculating, 363–364
- `snprintf()`
 - asynchronous signal safety, 515
 - copying overlapping objects, 81
 - return codes, checking, 432, 434–437
 - truncating null-terminated byte strings, 280–282
 - user input in format strings, 421
- Soft links, 617
- Software security, definition, xv–xvi
- Source character set, 590
- Source code, compliance validation, 3–4
- Source Code Analysis Tool Functional Specification*, 331
- Spaces, in file names, 591
- `spc_sanitize_environment()`, 479
- Special characters
 - ASCII, 590
 - national use positions, 590
 - portability, 590
 - sanitizing, 276–280
- Special file access, 426–431
- `sprintf()`
 - copying overlapping objects, 81
 - in-band errors, 537–538
 - truncating null-terminated byte strings, 280–282
- `sprintf_m()`, 538
- `sqrt()`, 229
- Square brackets ([]), subscript operator, 242
- `sscanf()`
 - converting strings to integers, 160–161
 - copying overlapping objects, 81
 - end-of-file indicator, 440
- Stack allocation, 335–338
- Stack exhaustion, 255
- Static allocation, 275–276
- Static assertions, 19, 39–42
- Static functions, macro alternative, 6–11, 30–32
- static objects, 70–72
- `static_assert()` macro, 41–42
- `stddef.h`, 135
- `stdint.h`, 142–143
- `stdio.h`, 106
- `stdlib.h`, 607
- Storage allocation, null-terminated
 - byte strings, 294–299
- Stoughton, Nick, 341
- `strcat()`, 280–282, 288–291
- `strcat_s()`, 288–291
- `strchr()`, 133, 441–442
- `strcoll()`, 560
- `strcpy()`
 - copying overlapping objects, 81
 - null-terminated byte strings
 - copying data, 296–297
 - memory allocation, 296–297
 - truncating, 280–282
 - replacement for, 288–291
- `strcpy_s()`, 281–282, 288–291
- `strdup()`, 471–473, 488–489
- Stream alteration return codes, checking, 408–409
- `strerror()`, 560
- String literals. *See* Characters and strings, string literals.
- Stringification, macro parameters, 12, 18–20
- Strings. *See* Characters and strings.
- `strlen()`
 - environment variables
 - copying, 298–299
 - size, calculating, 474–475
 - string length, determining, 296
 - wide character string size, determining, 304
- `strncat()`, 280–282, 288–291
- `strncat_s()`, 288–291

- `strncpy()`
 copying overlapping objects, 81
 null-terminating byte strings,
 300–301
 replacement for, 288–291
 truncating null-terminated byte
 strings, 280–282
- `strncpy_s()`, 288–291, 301
- `strtod()`, 559
- `strtok()`, 286–287
- `strtol()`, 158–162
- `strtol1()`, 159–162
- `strtol1()`, 159–162, 560–561
- `strtoul1()`, 159–162
- struct member alignment, 98–100
- Structured exception handling
 (SEH), 222–223
- Structures
 aligning in memory, 100–102
 byte-by-byte comparison,
 100–102
 element offset, determining,
 135–137
 packed, 99
 padding, 98–100
 size, as sum of parts, 98–100
- `strxfrm()`, 560
- Subscope, reusing variable names, 8
- Subscripts, and pointers, 244
- Subtraction
 integer overflow, 194–195
 pointers, 261–263, 263–265,
 265–269
 unsigned integer wrapping,
 183–184
- Sun Solaris TELNET daemon
 vulnerability, 278
- Survivability, 533–534
- Svoboda, David, 23
- Swapping. *See* Paging.
- `switch` statement, 573–574
- Symbolic links, 617
- `syslog()`, 422–423
- `syslog_r()`, 515
- `system()`, 479–487
- System qualities, compliant systems,
 1–2
- T**
- Temporary files
 creating in shared directories,
 454–463
 criteria for, 455
- naming, 455, 459
 predictable names, 455, 459
 removing, 456–457
 vulnerabilities, 457–461
- Testing
 with assertions, 597–598
 compiler diagnostic messages,
 571
 constant expression values, 19,
 39–42
 functions for error conditions,
 57–59
- Text mode vs. binary mode,
 411–413, 442–443
- `tgamma()`, 229
- `tgmath.h`, 226–227
- 3Rs of survivability, 534
- Three-character sequences, 22–24
- Tilde (~), complement operator,
 174–175
- Time values, 580–582, 610–613
- `time_t()`, 580–582, 610–613
- `tmpfile()`, 455, 460–461
- `tmpfile_s()`
 creating temporary binary files,
 460–461
 exceptions for use, 462
 temporary file criteria, 455
- `tmpnam()`, 455, 457–459
- `tmpnam_s()`, 455, 458–459
- `toascii()`, 314
- TOCTOU race conditions. *See also*
 Race conditions.
 accessing devices as files, 429
 checking file names, 378, 382
 unintended file access, 374
- Token pasting, 19
- Tokens, concatenating, 12, 18–20,
 29–30
- `tolower()`, 314
- Tool selection, 3–4
- Tool-generated code, 2–3
- Tools and utilities, xxii
- `toupper()`, 314
- Translation units, global variable ini-
 tialization order, 69–70
- Traversing arrays. *See* Iterating;
 Loops.
- Trigraph sequences, 22–24
- Truncation
 buffer overflow cause, 190
`fgets()`, 280–282
`gets()`, 280–282
- improper use of `strtok()`,
 286–287
- new-line character, missing,
 440–442
- null-terminated byte strings,
 280–282, 300–301
- pushing back characters, 411
- `snprintf()`, 280–282
- `sprintf()`, 280–282
- `strcat()`, 280–282
- `strcpy()`, 280–282
- `strcpy_s()`, 281–282
- string tokenization, 286–287
- `strncat()`, 280–282
- `strncpy()`, 280–282
- toward zero, 169
- `ungetc()`, 411
- `ungetwc()`, 411
- Type. *See* Data types.
- `typedef`
 vs. `#define`, 15–16
 readability, 44–45
- Type-generic functions, 10–11
- U**
- `uintmax_t`, 178–181
- `umask()`, 395
- Unary negation, 198–199
- Unbounded copying, 307–311
- Undefined behavior, 604–607
- `ungetc()`
 end-of-file, detecting, 436
 line breaks, 412
 new-line character, 412
 pushing back characters, 409–411
 return codes, checking, 388
 truncation, 411
- `ungetwc()`, 411
- Unintentional file access, 374,
 383–386, 463
- Union, 86
- Unique file names
 criteria for temporary files, 455
 ensuring, 461–462
 header file names, 17, 24–26
 predictability, 455, 459
 vulnerabilities, 457–461
- Unique identifiers, 78–79
- Universal character names, creating,
 29–30
- `unlink()`, 399–401, 486
- Unlinking files, 486. *See also* Remov-
 ing files.

- UnlockFile(), 456
 - Unlocking mutexes, 625–626
 - Unsafe macros, 30–32
 - unsetenv(), 479–482
 - unsigned char type, 162–164
 - unsigned character type, 274, 305–307
 - Unsigned integer conversions, 187–189
 - Unsigned integer wrapping, 164–165, 181–186
 - Unspecified nonportability, 603
 - Update mode, 444–446
 - US-ASCII characters, 590–593
 - usefile(), 588–590
 - User input in format strings, 421–424
 - ~username vs. ~, 482
 - Usual arithmetic conversions, 150–151
 - UTF-8 character encoding, 594–597
- V**
- va_arg() macro, 61, 62–64
 - valid(), 350
 - Values, removing unused, 600–602
 - Variable length arrays, 254–255, 335–336
 - Variable names in subscopes, reusing, 8, 36–38
 - Variables
 - global initialization order, assumptions about, 69–70
 - multiple per declaration, 16, 42–44
 - readability, 42–44
 - reusing names in subscopes, 36–38
 - size, determining, 49, 109–111
 - uninitialized, referencing, 124–128
 - unintentional reference, 36–39
 - visually indistinct identifiers, 38–39
 - Variadic functions. *See* Functions, variadic.
 - Venema, Wietsse, 280
 - Vertical bar (|), bitwise inclusive OR operator, 174–175
 - Vertical bars (||), logical OR, 96–98, 113
 - vfork()
 - denial-of-service attacks, 629
 - deprecated use, 629–631
 - portability issues, 629
 - security issues, 629
 - vfprintf(), 61
 - vfscanf(), 61, 157–158, 440
 - Viega, John, 277, 397
 - Virtual drives, 617
 - VirtualLock(), 340
 - Visual distinction, identifiers, 38–39
 - Visual Studio, runtime constraint handlers, 542–543
 - void return type, 118–119
 - Volatile access side effects, 30–32
 - Volatile qualifications, casting away, 123–124
 - volatile-qualified type objects, 82–84
 - vprintf(), 61
 - vscanf(), 61, 157–158, 440
 - vsprintf()
 - argument list warnings, 61
 - asynchronous signal safety, 515
 - return codes, 432
 - uses for, 26–27
 - variadic functions, 61
 - vsscanf(), 61
 - Vulnerabilities. *See also specific vulnerabilities.*
 - avoiding, 534–535
 - CVSS (Common Vulnerability Scoring System), xxvi
 - date public, xxvii
 - external programs, 478–482
 - file names, 372–374
 - ID numbers, xxvii
 - metrics, xxvi
 - names, xxvii
 - non-unique file names, 457–461
 - permission escalation, 463, 516
 - predictable file names, 457–461
 - Related Vulnerabilities sections, xxv–xxvii
 - severity, measuring, xxvi
 - signal handlers, 516, 522
 - temporary files, 457–461
- W**
- W32.Blaster.Worm, 259–260
 - Warning levels, compiler, 570–572
 - Warren, Henry S., 155
 - wchar_t type, 274
 - wcschar(), 441
 - wcsco11(), 560
 - wcslen(), 304
 - wcstod(), 559
 - wcsto1(), 559
 - wcstombs(), 81
 - wcsxfrm(), 560
 - _wdupenv_s(), 470–471
 - Web site resources. *See* Online resources.
 - wgetenv_s(), 469–470
 - White listing, characters and strings, 277–278
 - Wide character strings
 - buffer overflows, 303–305
 - data type, 274
 - definition, 273
 - modifying, 293–294
 - sizing, 303–305
 - string literals, 285
 - Wiki, CERT Secure Coding Standards, xvii
 - Wrapping
 - integers. *See* Unsigned integer wrapping.
 - memory allocation, checking for, 342–343
 - Writer/caller contract, 59–62
 - Writing data
 - binary, 401–403
 - to disk, 338–341
 - overwriting data, 473
- X**
- X Window System server vulnerability, 575–576
- Z**
- Zero-length space allocation, 332–335
 - ZeroMemory(), 583