

Information Technology — Programming Languages — Guidance to Avoiding Vulnerabilities in Programming Languages through Language Selection and Use

Élément introductif — Élément principal — Partie n: Titre de la partie

Warning

This document is not an ISO International Standard. It is distributed for review and comment. It is subject to change without notice and may not be referred to as an International Standard.

Recipients of this draft are invited to submit, with their comments, notification of any relevant patent rights of which they are aware and to provide supporting documentation.

Copyright notice

This ISO document is a working draft or committee draft and is copyright-protected by ISO. While the reproduction of working drafts or committee drafts in any form for use by participants in the ISO standards development process is permitted without prior permission from ISO, neither this document nor any extract from it may be reproduced, stored or transmitted in any form for any other purpose without prior written permission from ISO.

Requests for permission to reproduce this document for the purpose of selling it should be addressed as shown below or to ISO's member body in the country of the requester:

*ISO copyright office
Case postale 56, CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org*

Reproduction for sales purposes may be subject to royalty payments or a licensing agreement.

Violators may be prosecuted.

Contents

Page

Foreword	viii
Introduction	ix
1 Scope	1
1.1 In Scope	1
1.2 Not in Scope	1
1.3 Approach	1
1.4 Intended Audience	1
1.5 How to Use This Document	2
2 Normative references	5
3 Terms and definitions	6
4 Symbols (and abbreviated terms)	8
5 Vulnerability issues	9
5.1 Issues arising from incomplete or evolving language specifications	9
5.2 Issues arising from human cognitive limitations	11
5.3 Issues arising from a lack of predictable execution	12
5.4 Issues arising from the lack of portability and interoperability	12
5.5 Issues arising from inadequate language intrinsic support	13
5.6 Issues arising from language features prone to erroneous use	13
6 Programming Language Vulnerabilities	14
6.1 General	14
6.2 Obscure Language Features [BRS]	14
6.3 Unspecified Behaviour [BQF]	15
6.4 Undefined Behaviour [EWF]	16
6.5 Implementation-defined Behaviour [FAB]	17
6.6 Deprecated Language Features [MEM]	19
6.7 Pre-processor Directives [NMP]	20
6.8 Choice of Clear Names [NAI]	21
6.9 Choice of Filenames and other External Identifiers [AJN]	23
6.10 Unused Variable [XYR]	24
6.11 Identifier Name Reuse [YOW]	25
6.12 Namespace Issues [BJL]	27
6.13 Type System [IHN]	29
6.14 Bit Representations [STR]	31
6.15 Floating-point Arithmetic [PLF]	32
6.16 Enumerator Issues [CCB]	34
6.17 Numeric Conversion Errors [FLC]	35
6.18 String Termination [CJM]	37
6.19 Boundary Beginning Violation [XYX]	38
6.20 Unchecked Array Indexing [XYZ]	39
6.21 Unchecked Array Copying [XYW]	41
6.22 Buffer Overflow [XZB]	42
6.23 Pointer Casting and Pointer Type Changes [HFC]	43
6.24 Pointer Arithmetic [RVG]	44
6.25 Null Pointer Dereference [XYH]	45
6.26 Dangling Reference to Heap [XYK]	45
6.27 Templates and Generics [SYM]	47
6.28 Inheritance [RIP]	49
6.29 Initialization of Variables [LAV]	50

6.30	Wrap-around Error [XYY]	52
6.31	Sign Extension Error [XZI]	53
6.32	Operator Precedence/Order of Evaluation [JCW].....	54
6.33	Side-effects and Order of Evaluation [SAM]	55
6.34	Likely Incorrect Expression [KOA].....	56
6.35	Dead and Deactivated Code [XYQ].....	58
6.36	Switch Statements and Static Analysis [CLL]	59
6.37	Demarcation of Control Flow [EOJ]	60
6.38	Loop Control Variables [TEX]	61
6.39	Off-by-one Error [XZH]	62
6.40	Structured Programming [EWD].....	64
6.41	Passing Parameters and Return Values [CSJ]	65
6.42	Dangling References to Stack Frames [DCM].....	67
6.43	Subprogram Signature Mismatch [OTR]	68
6.44	Recursion [GDL].....	69
6.45	Returning Error Status [NZN]	71
6.46	Termination Strategy [REU].....	73
6.47	Extra Intrinsics [LRM].....	74
6.48	Type-breaking Reinterpretation of Data [AMV].....	75
6.49	Memory Leak [XYL].....	76
6.50	Argument Passing to Library Functions [TRJ]	78
6.51	Dynamically-linked Code and Self-modifying Code [NYY].....	79
6.52	Library Signature [NSQ]	80
6.53	Unanticipated Exceptions from Library Routines [HJW].....	81
7.	Application Vulnerabilities.....	83
7.1	Adherence to Least Privilege [XYN].....	83
7.2	Privilege Sandbox Issues [XYO].....	83
7.3	Executing or Loading Untrusted Code [XYS]	84
7.4	Unspecified Functionality [BVQ].....	85
7.5	Distinguished Values in Data Types [KLG]	86
7.6	Memory Locking [XZX]	87
7.7	Resource Exhaustion [XZP].....	88
7.8	Injection [RST].....	89
7.9	Cross-site Scripting [XYT]	91
7.10	Unquoted Search Path or Element [XZQ].....	93
7.11	Improperly Verified Signature [XZR].....	94
7.12	Discrepancy Information Leak [XZL]	94
7.13	Sensitive Information Uncleared Before Release [XZK]	95
7.14	Path Traversal [EWR]	96
7.15	Missing Required Cryptographic Step [XZS]	98
7.16	Insufficiently Protected Credentials [XYM]	98
7.17	Missing or Inconsistent Access Control [XZN].....	99
7.18	Authentication Logic Error [XZO].....	99
7.19	Hard-coded Password [XYP]	101
A.	Guideline Selection Process.....	102
A.1	Cost/Benefit Analysis	102
A.2	Documenting of the selection process.....	102
B.	Skeleton template for use in proposing programming language vulnerabilities.....	103
B.1	6.<x> <short title> [<unique immutable identifier>]	103
C.	Skeleton template for use in proposing application vulnerabilities.....	105
C.1	7.<x> <short title> [<unique immutable identifier>]	105
D.	Vulnerability Outline	106
E.	Vulnerability descriptions for <language>	108
E.1	<language>.1 Identification of standards.....	108
E.2	<language>.2 General terminology and concepts.....	108
E.3	<language>.<x> <Vulnerability Name> [<3 letter tag>]	108

Bibliography.....110

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of technical committees is to prepare International Standards. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights.

ISO/IEC TR 24772, which is a Technical Report of type 3, was prepared by Joint Technical Committee ISO/IEC JTC 1, Subcommittee SC 22, Programming Languages.

Introduction

All programming languages contain constructs that are incompletely specified, exhibit undefined behaviour, are implementation-dependent, or are difficult to use correctly. The use of those constructs may therefore give rise to *vulnerabilities*, as a result of which, software programs can execute differently than intended by the writer. In some cases, these vulnerabilities can compromise the safety of a system or be exploited by attackers to compromise the security or privacy of a system.

This Technical Report is intended to provide guidance spanning multiple programming languages, so that application developers will be better able to avoid the programming constructs that lead to vulnerabilities in software written in their chosen language and their attendant consequences. This guidance can also be used by developers to select source code evaluation tools that can discover and eliminate some constructs that could lead to vulnerabilities in their software or to select a programming language that avoids anticipated problems.

It should be noted that this Technical Report is inherently incomplete. It is not possible to provide a complete list of programming language vulnerabilities because new weaknesses are discovered continually. Any such report can only describe those that have been found, characterized, and determined to have sufficient probability and consequence.

Furthermore, to focus its limited resources, the working group developing this report decided to defer comprehensive treatment of several subject areas until future editions of the report. These subject areas include:

- Object-oriented language features (Although some simple issues related to inheritance are described in RIP)
- Concurrency
- Numerical analysis (although some simple items regarding the use of floating point are described in PLF)
- Scripting languages
- Inter-language operability
- Language-specific annexes

1 Information Technology — Programming Languages — Guidance to Avoiding Vulnerabilities in Programming
2 Languages through Language Selection and Use

3 **1 Scope**

4 **1.1 In Scope**

5 This Technical Report specifies software programming language vulnerabilities that should be avoided in the
6 development of systems where assured behaviour is required for security, safety, mission critical and business
7 critical software. In general, this guidance is applicable to the software developed, reviewed, or maintained for any
8 application.

9 Vulnerabilities are described in a generic manner that is applicable to a broad range of programming languages.

10 **1.2 Not in Scope**

11 This Technical Report does not address software engineering and management issues such as how to design and
12 implement programs, use configuration management tools, use managerial processes, and perform process
13 improvement. Furthermore, the specification of properties to be assured is not treated.

14 The specification of an application is *not* within the scope.

15 While this document does not discuss specification or design issues, there is recognition that boundaries among
16 the various activities are not clear-cut. This document seeks to avoid the debate about where low-level design ends
17 and implementation begins by treating selected issues that some might consider design issues rather than coding
18 issues.

19 **1.3 Approach**

20 Guidelines based on this Technical Report are likely to be highly leveraged in that they are likely to affect many
21 times more people than the number that worked on them. Therefore guidelines such as these have the potential to
22 make large savings, for a small cost, or to generate large unnecessary costs, for little benefit. For this reasons, the
23 writers of this Technical Report have taken a cautious approach to identifying programming language
24 vulnerabilities. New vulnerability descriptions can be added over time, as experience and evidence are
25 accumulated.

26 **1.4 Intended Audience**

27 The intended audience for this Technical Report is those who are concerned with assuring the existence of a
28 critical property in the software of their system; that is, those who are developing, qualifying, or maintaining a
29 software system and need to avoid language constructs that could cause the software to execute in a manner other
30 than intended.

31 As described in the following paragraphs, developers of applications that have clear safety, security or mission
32 criticality are expected to be aware of the risks associated with their code and could use this document to ensure
33 that their development practices address the issues presented by the chosen programming languages, for example
34 by subsetting or providing coding guidelines.

35 That should not be taken to mean that other developers can ignore this document. A weakness in an application
36 that of itself has no direct criticality may provide the route by which an attacker gains control of a system or may
37 otherwise disrupt co-hosted applications that are safety, security or mission critical.

1 It is hoped that such developers would use this document to ensure that common vulnerabilities are removed or at
2 least minimized from all applications.

3 Several specific audiences for this International Technical Report have been identified and are described below.

4 **1.4.1 Safety-Critical Applications**

5 Users who may benefit from this document include those developing, qualifying, or maintaining a system where it is
6 critical to prevent behaviour that might lead to:

- 7 • loss of human life or human injury, or
- 8 • damage to the environment.

9 **1.4.2 Security-Critical Applications**

10 Users who may benefit from this document include those developing, qualifying, or maintaining a system where it is
11 critical to exhibit security properties of:

- 12 • confidentiality,
- 13 • integrity, and
- 14 • availability.

16 **1.4.3 Mission-Critical Applications**

17 Users who may benefit from this document include those developing, qualifying, or maintaining a system where it is
18 critical to prevent behaviour that might lead to:

- 19 • property loss or damage, or
- 20 • economic loss or damage.

22 **1.4.4 Business-Critical Systems**

23 Users who may benefit from this document include those developing, qualifying, or maintaining a system whose
24 correctness and integrity are essential for successful operation of a business or enterprise.

25 **1.4.5 Modeling and Simulation Applications**

26 People who may benefit from this document include those who are primarily experts in areas other than
27 programming but need to use computation as part of their work. Such people include scientists, engineers,
28 economists, and statisticians. They require high confidence in the applications they write and use because of the
29 increasing complexity of the calculations made (and the consequent use of teams of programmers each
30 contributing expertise in a portion of the calculation), or to the costs of invalid results, or to the expense of individual
31 calculations implied by a very large number of processors used and/or very long execution times needed to
32 complete the calculations. These circumstances give a consequent need for high reliability and motivate the need
33 felt by these programmers for the guidance offered in this document.

34 **1.5 How to Use This Document**

35 This Technical Report gathers language-independent descriptions of programming language vulnerabilities, as well
36 as selected application vulnerabilities, which have occurred in the past and are likely to occur again. Because new
37 vulnerabilities are always being discovered, it is anticipated that the document will be revised and new descriptions
38 added. For that reason, a scheme that is distinct from document sub-clause numbering has been adopted to
39 identify the vulnerability descriptions. Each description has been assigned an arbitrarily generated, unique three-
40 letter code. These codes should be used in preference to sub-clause numbers when referencing descriptions.

1 The main part of the document contains descriptions that are intended to be language-independent to the greatest
2 possible extent. Future editions will include annexes that apply the generic guidance to particular programming
3 languages.

4 The document has been written with several possible usages in mind:

- 5 • Programmers familiar with the vulnerabilities of a specific language can reference the guide for more
6 generic descriptions and their manifestations in less familiar languages.
- 7 • Tool vendors can use the three-letter codes as a succinct way to “profile” the selection of vulnerabilities
8 considered by their tools.
- 9 • Individual organizations may wish to write their own coding standards intended to reduce the number of
10 vulnerabilities in their software products. The guide can assist in the selection of vulnerabilities to be
11 addressed in those standards and the selection of coding guidelines to be enforced.
- 12 • Organizations or individuals selecting a language for use in a project may want to consider the
13 vulnerabilities inherent in various candidate languages.

14 The following clauses include suggestions for ways of avoiding the vulnerabilities. It should be noted that these
15 include techniques that can be applied during development, and those that must be implemented as run-time
16 checks. The former are likely to be appropriate to all applications. For some applications, it is relatively more
17 important to ensure that potential run-time errors are eliminated during development because there may be
18 insufficient opportunity to recover from them. For long-running simulations, run-time checks may increase the run-
19 time to the point where the prediction loses value, or to the point where the cost of calculation becomes prohibitive
20 given the value of the simulation results. Source code checking tools can be used to automatically enforce some
21 coding rules and standards.

22 Sections 2, 3 and 4 provide Normative References, Terms and Definitions and symbols and definitions (including
23 acronyms), respectively.

24 Section 5, *Vulnerability Issues*, provides rationale for the document and explains how many of the vulnerabilities
25 occur.

26 Section 6, *Programming Language Vulnerabilities*, provides language-independent descriptions of vulnerabilities in
27 programming languages that can lead to application vulnerabilities. Each description provides:

- 28 • a summary of the vulnerability,
- 29 • characteristics of languages where the vulnerability may be found,
- 30 • typical mechanisms of failure,
- 31 • techniques that programmers can use to avoid the vulnerability, and
- 32 • ways that language designers can modify language specifications in the future to help programmers
33 mitigate the vulnerability.

34 Section 7, *Application Vulnerabilities*, provides descriptions of selected application vulnerabilities which have been
35 found and exploited in a number of applications and which have well known mitigation techniques, and which result
36 from design decisions made by coders in the absence of suitable language library routines or other mechanisms.
37 For these vulnerabilities, each description provides:

- 38 • a summary of the vulnerability,
- 39 • typical mechanisms of failure, and
- 40 • techniques that programmers can use to avoid the vulnerability.

41 Annexes A through C are templates and guidelines that were used in the identification, selection, and creation of
42 the vulnerabilities that were generated for sections 6 and 7. They can be used to guide the generation of future
43 language vulnerabilities and application vulnerabilities.

- 1 Annex D, *Vulnerability Outline*, is a categorization of the vulnerabilities of this report in the form of a hierarchical
- 2 outline.

- 3 Annex E, *Language Specific Vulnerability Template*, is a template for the writing of programming language specific
- 4 annexes that explain how the vulnerabilities from section 6 are realized in that programming language (or show
- 5 how they are absent), and how they might be mitigated in language-specific terms. Future revisions of this
- 6 Technical Report are planned to contain language-specific annexes that are developed using Annex E.

1 **2 Normative references**

2 The following referenced documents are indispensable for the application of this document. For dated references,
3 only the edition cited applies. For undated references, the latest edition of the referenced document (including any
4 amendments) applies.

5 ISO/IEC 9899:1999, *Programming Languages – C*

6 ISO/IEC 1539-1:2004, *Programming Languages – Fortran*

7 ISO/IEC 8652:1995/Cor 1:2001/Amd 1:2007, *Programming languages – Ada*

8 ISO/IEC 14882:2003, *Programming Languages – C++*

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply. Other terms are defined where they appear in *italic* type. Programming language keywords or syntactic tokens appear in `courier` font.

3.1 Language Vulnerability

A *property* (of a programming language) that can contribute to, or that is strongly correlated with, application vulnerabilities in programs written in that language.

Note: The term "property" can mean the presence or the absence of a specific feature, used singly or in combination. As an example of the absence of a feature, encapsulation (control of where names may be referenced from) is generally considered beneficial since it narrows the interface between modules and can help prevent data corruption. The absence of encapsulation from a programming language can thus be regarded as a vulnerability. Note that a property together with its complement may both be considered language vulnerabilities. For example, automatic storage reclamation (garbage collection) can be a vulnerability since it can interfere with time predictability and result in a safety hazard. On the other hand, the absence of automatic storage reclamation can also be a vulnerability since programmers can mistakenly free storage prematurely, resulting in dangling references.

3.2 Application Vulnerability

A security vulnerability or safety hazard, or defect.

3.3 Security Vulnerability

A weakness in an information system, system security procedures, internal controls, or implementation that could be exploited or triggered by a threat.

3.4 Safety Hazard

IEC61508 part 4: defines a "Hazard" as a "potential source of harm", where "harm" is "physical injury or damage to the health of people either directly or indirectly as a result of damage to property or to the environment".

IEC61508 cites ISO/IEC Guide 51 as the source for the definition.

Note: IEC61508 is titled "Functional safety of electrical/electronic/ programmable electronic safety-related systems", with part 4 being "Definitions and abbreviations". Hence within IEC61508 the "safety" context of "safety hazard" is assumed.

Note: Some derived standards, such as UK Defence Standard 00-56, broaden the definition of "harm" to include material and environmental damage (not just harm to people caused by property and environmental damage).

3.5 Safety-critical software

Software for applications where failure can cause very serious consequences such as human injury or death. IEC61508 part 4: defines "Safety-related software" as "software that is used to implement safety functions in a safety-related system.

Note: Notwithstanding that in some domains a distinction is made between safety-related (may lead to any harm) and safety-critical (life threatening), this Technical Report uses the term *safety-critical* for all vulnerabilities that may result in safety hazards.

3.6 Software quality

1 The degree to which software implements the requirements described by its specification and the degree to which
2 the characteristics of a software product fulfils its requirements.

3 3.7 Predictable Execution

4 The property of the program such that all possible executions have results that can be predicted from the source
5 code, the relevant language-defined implementation characteristics and knowledge of the universe of execution.

6 **Note:** In some environments, this would raise issues regarding numerical stability, exceptional processing, and
7 concurrent execution.

8 **Note:** Predictable execution is an ideal that must be approached keeping in mind the limits of human
9 capability, knowledge, availability of tools, and other factors. Neither this nor any standard ensures predictable
10 execution. Rather this standard provides advice on improving predictability. The purpose of this document is to
11 assist a reasonably competent programmer to approach the ideal of predictable execution.

12 **Note:** The following terms are used in relation to “Predictable execution”

- 13 • **Unspecified behaviour:** A situation where the implementation of a language will have to make some
14 choice from a finite set of alternatives, but that choice is not in general predictable by the programmer, for
15 example, the order in which sub-expressions are evaluated in an expression in many languages.
- 16 • **Implementation-defined behaviour:** A situation where the implementation of a language will have to
17 make some choice, and it is required that this choice be documented and available to the programmer.
- 18 • **Undefined behaviour:** A situation where the definition of a language can give no indication of what
19 behaviour to expect from a program – it may be some form of catastrophic failure (a ‘crash’) or continued
20 execution with some arbitrary data.

21 **Note:** This document includes a section on **Unspecified functionality**. This notion is related to neither
22 unspecified behaviour, which is a characteristic of an application, nor the language used to develop the
23 application.

24

1 **4 Symbols (and abbreviated terms)**

2 None.

5 Vulnerability issues

Software vulnerabilities are unwanted characteristics of software that may allow software to behave in ways that are unexpected by a reasonably sophisticated user of the software. The expectations of a reasonably sophisticated user of software may be set by the software's documentation or by experience with similar software. Programmers introduce vulnerabilities into software by failing to understand the expected behaviour (the software requirements), or by failing to correctly translate the expected behaviour into the actual behaviour of the software.

This document does not discuss a programmer's understanding of software requirements or the completeness of the requirements. This document does not discuss software engineering issues per se. This document does not discuss configuration management, build environments, code-checking tools, nor software testing. This document does not discuss the classification of software vulnerabilities according to safety or security concerns. This document does not discuss the costs of software vulnerabilities, or the costs of preventing them.

This document does discuss a reasonably competent programmer's failure to translate the understood requirements into correctly functioning software. This document does discuss programming language features known to contribute to software vulnerabilities. That is, this document discusses issues arising from those features of programming languages found to increase the frequency of occurrence of software vulnerabilities. The intention is to provide guidance to those who wish to specify coding guidelines for their own particular use.

A programmer writes source code in a programming language to translate the understood requirements into working software. The programmer selects and codes constructs specified by a programming language with the intention of achieving a written expression of the desired behaviour.

A program's expected behaviour might be stated in a complex technical document, which can result in a complex sequence of features of the programming language. Software vulnerabilities occur when a reasonably competent programmer fails to understand the totality of the effects of the language features combined to construct the software. The overall software may be a very complex technical document itself (written in a programming language whose definition is also a complex technical document).

The recommendations contained in this Technical Report might also be considered to be code quality issues. Both kinds of issues might be addressed through the use of a systematic development process, use of development/analysis tools and thorough testing.

5.1 Issues arising from incomplete or evolving language specifications

While there are many millions of programmers in the world, there are only several hundreds of authors engaged in designing and specifying those programming languages defined by international standards. The design and specification of a programming language is very different from programming. Programming involves selecting and sequentially combining features from the programming language to (locally) implement specific steps of the software's design. In contrast, the design and specification of a programming language involves (global) consideration of all aspects of the programming language. This must include how all the features will interact with each other, and what effects each will have, separately and in any combination, under all foreseeable circumstances. Thus, language design has global elements that are not generally present in any local programming task.

The creation of the abstractions which become programming language standards therefore involve consideration of issues unneeded in many cases of actual programming. Therefore perhaps these issues are not routinely considered when programming in the resulting language. These global issues may motivate the definition of subtle distinctions or changes of state not apparent in the usual case wherein a particular language feature is used. Authors of programming languages may also desire to maintain compatibility with older versions of their language while adding more modern features to their language and so add what appears to be an inconsistency to the language.

For example, some languages may allow a subprogram to be invoked without specifying the correct signature of the subprogram. This may be allowed to keep compatibility with earlier versions of the language where such usage was permitted, and despite the knowledge that modern practice demands the signature be specified. Specifically,

1 the programming language C does not require a function prototype be within scope¹. The programming language
 2 Fortran does not require an explicit interface. Thus, language usage is improved by coding standards specifying
 3 that the signature be present.

4 A reasonably competent programmer therefore may not consider the full meaning of every language feature used,
 5 as only the desired (local or subset) meaning may correspond to the programmer's immediate intention. In
 6 consequence, a subset meaning of any feature may be prominent in the programmer's overall experience.

7 Further, the combination of features indicated by a complex programming goal can raise the combinations of
 8 effects, making a complex aggregation within which some of the effects are not intended.

9 **5.1.1 Compiler Selection**

10 Compiler selection is important to ensure a system operates safely and securely. Compilers are important as they
 11 are the intermediary between the human readable source code and the machine readable binary code. This crucial
 12 step is often overlooked and compilers, unless coming from a trusted source and developed according to agreed
 13 standards, should be treated as any other commercial off the shelf software that has an unknown pedigree.

14 Often, developers analyze the source code to detect any code that can negatively impact security or safety. This
 15 aims to solve one part of the problem. After the source has been compiled, there is a need to be sure that the
 16 compiler did not insert any logic (maliciously or inadvertently) into the binary that compromises the systems security
 17 or safety. This is especially important because this type of vulnerability will be inserted into every piece of software
 18 that the compiler processes.

19 To combat against this, developers of security or safety critical systems should only use compilers from a trusted
 20 source. The trusted source should also provide evidence that the compiler is free from anomalous behaviour;
 21 similar to the way RTCA's DO-178B defines qualifiable tools. In addition, developers of critical software can
 22 perform source to binary traceability to ensure the compiler has not inserted any undesired logic into the binary
 23 code.

24 If a compiler has many options, then developers should consider if the options to be used in the project are the
 25 same options used when the compiler was validated.

26 **5.1.2 Issues arising from unspecified behaviour**

27 While every language standard attempts to specify how software written in the language will behave in all
 28 circumstances, there will always be some behaviour that is not specified completely. In any circumstance, of
 29 course, a particular compiler will produce a program with some specific behaviour (or fail to compile the program at
 30 all). Where a programming language construct is insufficiently defined, different compilers may generate different
 31 behaviours from the same source code. The authors of language standards often have an interpretations or
 32 defects process in place to treat these situations once they become known, and, eventually, to specify one
 33 particular behaviour. However, the time needed by the process to produce corrections to the language standard is
 34 often long, as careful consideration of the issues involved is needed.

35 When programs are compiled with only one compiler, the programmer may not be aware when behaviour not
 36 specified by the standard has been produced. Programs relying upon behaviour not specified by the language
 37 standard may behave differently when they are compiled with different compilers. An experienced programmer
 38 may choose to use more than one compiler, even in one environment, to obtain diagnostics from more than one
 39 source. In this usage, any particular compiler must be considered to be a different compiler if it is used with
 40 different options (which can give it different behaviour), or is a different release of the same compiler (which may
 41 have different default options or may generate different code), or is on different hardware (which may have a
 42 different instruction set). In this usage, a different computer may be the same hardware with a different operating
 43 system, with different compilers installed, with different software libraries available, with a different release of the
 44 same operating system, or with a different operating system configuration.

¹ This feature has been deprecated in the 1999 version of the ISO C Standard.

1 5.1.3 Issues arising from implementation-defined behaviour

2 In some situations, a programming language standard may specifically allow compilers to support a range of
3 behaviours to a given language feature or combination of features. This may enable a more efficient execution on a
4 wider range of hardware, or enable use of the programming language in a wider variety of circumstances.

5 To allow use on a wide range of hardware, for example, many languages do not specify the amount of storage
6 reserved for language-defined entities such as variables. The degree to which a diligent programmer may obtain
7 information on the amount of storage reserved for entities varies among languages.

8 The authors of language standards are encouraged to provide lists of all allowed variations of behaviour (as many
9 already do). Such a summary will benefit applications programmers, those who define applications coding
10 standards, and those who make code analysis tools.

11 5.1.4 Issues arising from undefined behaviour

12 In some situations, a programming language standard may specify that program behaviour is undefined. While the
13 authors of language standards naturally try to minimize these situations, they may be inevitable when attempting to
14 define software recovery from errors, or other situations recognized as being incapable of precise definition.

15 An example of undefined behaviour, in many languages, is the use of the value of a variable to which there has not
16 yet been an assignment.

17 5.2 Issues arising from human cognitive limitations

18 The goal in the creation of a programming language is to provide a tool that may be used by the writers of software
19 to manipulate data and produce a desired result. Some programming languages are general purpose languages,
20 while others are targeted to specific tasks or needs. Even general purpose languages may be created with a
21 specific user in mind, for example C was created by engineers for programmers implementing system software,
22 while COBOL was created for business analysts and business programmers.

23 All humans are different. Constructs that may be easily understood by mathematicians may be confusing to a
24 business analyst. Also, similar constructs in different languages (the use of '=' in COBOL for comparison, and in C
25 for assignment), and similar constructs within the same language (the use of parentheses for both function
26 parameter lists, macro parameter lists, and for array subscripts) can add to this confusion. In addition, many
27 languages provide multiple syntaxes to accomplish the same task, and coders will choose those syntaxes that
28 make the most sense to them, again adding additional confusion in program creation and maintenance.

29 Humans are also fallible, and can only comprehend a limited number of interactions within a process before that
30 process must be subdivided into smaller segments. In addition, stress, whether internal pressures and deadlines
31 or external influences totally unrelated to the task at hand, can interfere with the process of software creation.
32 These factors combined with language constructs that may be confusing can lead to failures due to human
33 fallibility. These failures can include:

- 34 • Cognitive shortcomings, which are failures of design and implementation. These may be
 - 35 ○ faulty reasoning and
 - 36 ○ incomplete solutions due to lack of time and effort in comprehending the task.
- 37 • Knowledge shortcomings, which are failures of training and environment. These can include
 - 38 ○ incomplete and/or incorrect knowledge of appropriate language semantics,
 - 39 ○ incomplete and/or incorrect knowledge of how the chosen syntax will be executed by a particular
40 implementation, and
 - 41 ○ incomplete and/or incorrect knowledge of the internal and external interactions of the various
42 software components involved.
- 43 • Judgment shortcomings, which are failures of choice and will, and may include
 - 44 ○ selection of simpler but vulnerable constructs in place of more robust but more time consuming
45 solutions and

- selection of terse constructs, which may be less understandable, in place of verbose, maintainable solutions.

This technical report identifies issues related to the use of programming languages that can increase the likelihood of errors due to cognitive limitations, and recommends ways that can be used to mitigate or eliminate these errors. Some of the mechanisms recommended in this technical report include reducing the cognitive effort necessary in reading existing source code, reducing the amount of knowledge needed by readers of existing source code, and reducing the probability that incorrect developer knowledge will result in unpredictable code execution.

5.3 Issues arising from a lack of predictable execution

If a reasonably competent programmer has a good understanding of the intent of a program after reading source code as far as a particular line of code, the programmer ought to have a good understanding of the state of the program after reading the next line of code. However, some features, or, more likely, some combinations of features, of programming languages are associated with relatively decreased rates of the programmer's maintaining their understanding as they read through a program. It is these features and combinations of features that are indicated in this document, along with ways to increase the programmer's understanding as code is read.

Here, the term understanding means the programmer's recognition of all effects, including subtle or unintended changes of state, of any language feature or combination of features appearing in the program. This view does not imply that programmers only read code from beginning to end. It is simply a statement that a line of code changes the state of a program, and that a reasonably competent programmer ought to understand the state of the program both before and after reading any line of code. It is only to a first approximation that code is read and understood line by line.

5.4 Issues arising from the lack of portability and interoperability

The representation of characters, the representation of true/false values, the set of valid addresses, the properties and limitations of any (fixed-point or floating-point) numerical quantities, and the representation of programmer-defined types and classes may vary among hardware, among languages (affecting inter-language software development), and among compilers of a given language. These variations may be the result of hardware differences, operating system differences, library differences, compiler differences, or different configurations of the same compiler (as may be set by environment variables or configuration files). In each of these circumstances, there is an additional burden on the programmer because part of the program's behaviour is indicated by a factor that is not a part of the source code. That is, the program's behaviour may be indicated by a factor that is invisible when reading the source code. Compilation control schemes (IDE projects, make, and scripts) further complicate this situation by abstracting and manipulating the relevant variables (target platform, compiler options, libraries, and so forth).

Many compilers of standard-defined languages also support language features that are not specified by the language standard. These non-standard features are called extensions. For portability, the programmer must be aware of the language standard, and use only constructs with standard-defined semantics. The motivation to use extensions may include the desire for increased functionality within a particular environment, or increased efficiency on particular hardware. There are well-known software engineering techniques for minimizing the ill effects of extensions; these techniques should be a part of any coding standard where they are needed, and they should be employed whenever extensions are used. These issues are software engineering issues and are not further discussed in this document.

Some language standards define libraries that are available as a part of the language definition. Such libraries are an intrinsic part of the respective language and are called intrinsic libraries. There are also libraries defined by other sources; these are called non-intrinsic libraries.

The use of non-intrinsic libraries to broaden the software primitives available in a given development environment is a useful technique, allowing the use of trusted functionality directly in the program. Libraries may also allow the program to bind to capabilities provided by an environment. However, these advantages are potentially offset by any lack of skill on the part of the designer of the library (who may have designed subtle or undocumented changes of state into the library's behaviour), and implementer of the library (who may not have implemented the library identically on every platform), and even by the availability of the library on a new platform. The quality of the

1 documentation of a third-party library is another factor that may decrease the reliability of software using a library in
2 a particular situation by failing to describe clearly the library's full behaviour. If a library is missing on a new
3 platform, its functionality must be recreated to port any software depending upon the missing library. The re-
4 creation may be burdensome if the reason the library is missing is because the underlying capability for a particular
5 environment is missing. The pedigree of the library must also be assured, the library must come from a trusted
6 source without adulteration in transmission or storage.

7 Using a non-intrinsic library usually requires that options be set during compilation and linking phases, which
8 constitute a software behaviour specification beyond the source code. Again, these issues are software
9 engineering issues and are not further discussed in this document.

10 **5.5 Issues arising from inadequate language intrinsic support**

11 Many languages are created to facilitate programming within an application domain. Some languages are
12 specifically designed for programming of business applications, numerical computation or systems programming.
13 Problems can arise when, for example, a language being used to implement a real-time, multi-threaded system
14 lacks key features that are needed such as a way of enforcing mutual exclusion. Such features can be provided by
15 the programming environment in the form of libraries, but the definition of such libraries may be proprietary and
16 inclined to change in later releases. A vendor may even decide to withdraw support entirely for such a library. Also,
17 such a library may not be verified and validated to the same standard as the compiler and the application being
18 developed.

19 Some potential problems may be preventable through the use of stronger types, or the use of controls such as
20 array bounds checking or integer checking to avoid overflows. These stronger restraints on a language have a cost
21 both in performance and in the flexibility to perform certain operations. Language designers must strike a balance
22 between restraints in the language, performance and flexibility causing some languages to lean heavily toward one
23 or more extremes in pursuit of some language attributes. The intrinsic support provided by a language can help
24 considerably in avoiding vulnerabilities, but such support can cause the utility of programming within a particular
25 application domain to diminish.

26 **5.6 Issues arising from language features prone to erroneous use**

27 Certain language constructs are relatively simple and straightforward to use. Other ones are complex to use or
28 easily misused in a legal, but unintended way. Programmers may use floating point variables and pointers without
29 fully understanding the nuances of the data representation. Rarely needed constructs or constructs that can be
30 substituted for a series of simpler constructs can be used without a complete understanding of the full effects of the
31 constructs.

32 Syntactic language features that are not intolerant of common typo errors can produce some problems that are
33 notoriously difficult to find. One common example of this is that C permits an unintentional assignment to be
34 performed in a Boolean expression by the accidental use of a single "=" (assignment) instead of the intended "=="
35 test for equality. It then allows the resulting value to be treated as a Boolean.

6. Programming Language Vulnerabilities

6.1 General

The standard for a programming language provides definitions for that language's constructs. This Technical Report will in general use the terminology that is most natural to the description for each individual vulnerability, relying upon the individual standards for terminology details. In general, the reader should be aware that "method", "function", and "procedure" could denote similar constructs in different languages, as can "pointer" and "reference". Situations described as "undefined behaviour" in some languages are known as "unbounded behaviour" in others.

6.2 Obscure Language Features [BRS]

6.2.1 Description of application vulnerability

Every programming language has features that are obscure, difficult to understand or difficult to use correctly. The problem is compounded if a software design must be reviewed by people who may not be language experts, such as, hardware engineers, human-factors engineers, or safety officers. Even if the design and code are initially correct, maintainers of the software may not fully understand the intent. The consequences of the problem are more severe if the software is to be used in trusted applications, such as safety or mission critical ones.

Misunderstood language features or misunderstood code sequences can lead to application vulnerabilities in development or in maintenance.

6.2.2 Cross reference

JSF AV Rules: 84, 86, 88, and 97
MISRA C 2004: 3.2, 10.2, 13.1, 20.6-20.12, 12.10, and 17.5
MISRA C++ 2008: 0-2-1, 2-3-1, and 12-1-1
CERT C guidelines: FIO03-C, MSC05-C, MSC30-C, and MSC31-C.
ISO/IEC TR 15942:2000: 5.4.2, 5.6.2 and 5.9.3

6.2.3 Mechanism of failure

The use of obscure language features can lead to an application vulnerability in several ways:

- The original programmer may misunderstand the correct usage of the feature and could utilize it incorrectly in the design or code it incorrectly.
- Reviewers of the design and code may misunderstand the intent or the usage and overlook problems.
- Maintainers of the code cannot fully understand the intent or the usage and could introduce problems during maintenance.

6.2.4 Applicable language characteristics

This vulnerability description is intended to be applicable to any language.

6.2.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Individual programmers should avoid the use of language features that are obscure or difficult to use, especially in combination with other difficult language features. Organizations should adopt coding standards that discourage use of such features or show how to use them correctly.

- 1 • Organizations developing software with critically important requirements should adopt a mechanism to
- 2 monitor which language features are correlated with failures during the development process and during
- 3 deployment.
- 4 • Organizations should adopt or develop stereotypical idioms for the use of difficult language features, codify
- 5 them in organizational standards, and enforce them via review processes.
- 6 • Avoid the use of complicated features of a language.
- 7 • Avoid the use of rarely used constructs that could be difficult for entry-level maintenance personnel to
- 8 understand.
- 9 • Static analysis can be used to find incorrect usage of some language features.

10 It should be noted that consistency in coding is desirable for each of review and maintenance. Therefore, the
 11 desirability of the particular alternatives chosen for inclusion in a coding standard does not need to be empirically
 12 proven.

13 6.2.6 Implications for standardization

14 In future standardization activities, the following items should be considered:

- 15 • Language designers should consider removing or deprecating obscure, difficult to understand, or difficult to
- 16 use features.
- 17 • Language designers should provide language directives that optionally disable obscure language features.

18 6.3 Unspecified Behaviour [BQF]

19 6.3.1 Description of application vulnerability

20 The external behaviour of a program whose source code contains one or more instances of constructs having
 21 unspecified behaviour may not be fully predictable when the source code is (re)compiled or (re)linked.

22 6.3.2 Cross reference

23 JSF AV Rules: 17-25

24 MISRA C 2004: 1.3, 1.5, 3.1 3.3, 3.4, 17.3, 1.2, 5.1, 18.2, 19.2, and 19.14

25 MISRA C++ 2008: 5-0-1, 5-2-6, 7-2-1, and 16-3-1

26 CERT C guidelines: MSC15-C

27 See: Undefined Behaviour [EWF] and Implementation-defined Behaviour [FAB].

28 6.3.3 Mechanism of failure

29 Language specifications do not always uniquely define the behaviour of a construct. When an instance of a
 30 construct that is not uniquely defined is encountered (this might be at any of compile, link, or run time)
 31 implementations are permitted to choose from the set of behaviours allowed by the language specification. The
 32 term 'unspecified behaviour' is sometimes applied to such behaviours, (language specific guidelines need to
 33 analyze and document the terms used by their respective language).

34 A developer may use a construct in a way that depends on a subset of the possible behaviours occurring. The
 35 behaviour of a program containing such a usage is dependent on the translator used to build it always selecting the
 36 'expected' behaviour.

37 Many language constructs may have unspecified behaviour and unconditionally recommending against any use of
 38 these constructs may be impractical. For instance, in many languages the order of evaluation of the operands
 39 appearing on the left- and right-hand side of an assignment is unspecified, but in most cases the set of possible
 40 behaviours always produce the same result.

1 The appearance of unspecified behaviour in a language specification is a recognition by the language designers
2 that in some cases flexibility is needed by software developers and provides a worthwhile benefit for language
3 translators; this usage is not a defect in the language.

4 The important characteristic is not the internal behaviour exhibited by a construct (such as the sequence of
5 machine code generated by a translator) but its external behaviour (that is, the one visible to a user of a program).
6 If the set of possible unspecified behaviours permitted for a specific use of a construct all produce the same
7 external effect when the program containing them is executed, then rebuilding the program cannot result in a
8 change of behaviour for that specific usage of the construct.

9 For instance, while the following assignment statement contains unspecified behaviour in many languages (that is,
10 it is possible to evaluate either the A or B operand first, followed by the other operand):

11 `A = B;`

12 in most cases the order in which A and B are evaluated does not effect the external behaviour of a program
13 containing this statement.

14 **6.3.4 Applicable language characteristics**

15 This vulnerability is intended to be applicable to languages with the following characteristics:

- 16 • Languages whose specification allows a finite set of more than one behaviour for how a translator handles
17 some construct, where two or more of the behaviours can result in differences in external program
18 behaviour.

19 **6.3.5 Avoiding the vulnerability or mitigating its effects**

20 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 21 • Use language constructs that have specified behaviour.
- 22 • Ensure that a specific use of a construct having unspecified behaviour produces a result that is the same
23 for all of the possible behaviours permitted by the language specification.
- 24 • When developing coding guidelines for a specific language all constructs that have unspecified behaviour
25 should be documented and for each construct the situations where the set of possible behaviours can vary
26 should be enumerated.

27 **6.3.6 Implications for standardization**

28 In future standardization activities, the following items should be considered:

- 29 • Languages should minimize the amount of unspecified behaviours, minimize the number of possible
30 behaviours for any given "unspecified" choice, and document what might be the difference in external
31 effect associated with different choices.

32 **6.4 Undefined Behaviour [EWF]**

33 **6.4.1 Description of application vulnerability**

34 The external behaviour of a program containing an instance of a construct having undefined behaviour, as defined
35 by the language specification, is not predictable.

36 **6.4.2 Cross reference**

37 JSF AV Rules: 17-25

38 MISRA C 2004: 1.3, 1.5, 3.1 3.3, 3.4, 17.3, 1.2, 5.1, 18.2, 19.2, and 19.14

1 MISRA C++ 2008: 2-13-1, 5-2-2, 16-2-4, and 16-2-5
 2 CERT C guidelines: MSC15-C
 3 See: Unspecified Behaviour [BQF] and Implementation-defined Behaviour [FAB].

4 **6.4.3 Mechanism of failure**

5 Language specifications may categorize the behaviour of a language construct as undefined rather than as a
 6 semantic violation (that is, an erroneous use of the language) because of the potentially high implementation cost
 7 of detecting and diagnosing all occurrences of it. In this case no specific behaviour is required and the translator or
 8 runtime system is at liberty to do anything it pleases (which may include issuing a diagnostic).

9 The behaviour of a program built from successfully translated source code containing a construct having undefined
 10 behaviour is not predictable. For example, in some languages the value of a variable is undefined before it is
 11 initialized.

12 **6.4.4 Applicable language characteristics**

13 This vulnerability is intended to be applicable to languages with the following characteristics:

- 14 • Languages that do not fully define the extent to which the use of a particular construct is a violation of the
 15 language specification.
- 16 • Languages that do not fully define the behaviour of constructs during compile, link and program execution.

17 **6.4.5 Avoiding the vulnerability or mitigating its effects**

18 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 19 • Ensuring that undefined language constructs are not used.
- 20 • Ensuring that a use of a construct having undefined behaviour does not operate within the domain in which
 21 the behaviour is undefined. When it is not possible to completely verify the domain of operation during
 22 translation a runtime check may need to be performed.
- 23 • When developing coding guidelines for a specific language all constructs that have undefined behaviour
 24 should be documented. The items on this list might be classified by the extent to which the behaviour is
 25 likely to have some critical impact on the external behaviour of a program (the criticality may vary between
 26 different implementations, for example, whether conversion between object and function pointers has well
 27 defined behaviour).

28 **6.4.6 Implications for standardization**

29 In future standardization activities, the following items should be considered:

- 30 • Language designers should minimize the amount of undefined behaviour to the extent possible and
 31 practical.
- 32 • Language designers should enumerate all the cases of undefined behaviour.
- 33 • Language designers should provide mechanisms that permit the disabling or diagnosing of constructs that
 34 may produce undefined behaviour.

35 **6.5 Implementation-defined Behaviour [FAB]**

36 **6.5.1 Description of application vulnerability**

37 Some constructs in programming languages are not fully defined (see Unspecified Behaviour [BQF]) and thus
 38 leave compiler implementations to decide how the construct will operate. The behaviour of a program whose
 39 source code contains one or more instances of constructs having implementation-defined behaviour, can change
 40 when the source code is recompiled or relinked.

6.5.2 Cross reference

JSF AV Rules: 17-25
 MISRA C 2004: 1.3, 1.5, 3.1 3.3, 3.4, 17.3, 1.2, 5.1, 18.2, 19.2, and 19.14
 MISRA C++ 2008: 5-2-9, 5-3-3, 7-3-2, and 9-5-1
 CERT C guidelines: MSC15-C
 ISO/IEC TR 15942:2000: 5.9
 Ada Quality and Style Guide: 7.1.5 and 7.1.6
 See: Unspecified Behaviour [BQF] and Undefined Behaviour [EWF].

6.5.3 Mechanism of failure

Language specifications do not always uniquely define the behaviour of a construct. When an instance of a construct that is not uniquely defined is encountered (this might be at any of translation, link-time, or program execution) implementations are permitted to choose from a set of behaviours. The only difference from unspecified behaviour is that implementations are required to document how they behave.

A developer may use a construct in a way that depends on a particular implementation-defined behaviour occurring. The behaviour of a program containing such a usage is dependent on the translator used to build it always selecting the 'expected' behaviour.

Some implementations provide a mechanism for changing an implementation's implementation-defined behaviour (for example, use of `pragmas` in source code). Use of such a change mechanism creates the potential for additional human error in that a developer may be unaware that a change of behaviour was requested earlier in the source code and may write code that depends on the implementation-defined behaviour that occurred prior to that explicit change of behaviour.

Many language constructs may have implementation-defined behaviour and unconditionally recommending against any use of these constructs may be completely impractical. For instance, in many languages the number of significant characters in an identifier is implementation-defined. Developers need to choose a minimum number of characters and require that only translators supporting at least that number, N , of characters be used.

The appearance of implementation-defined behaviour in a language specification is recognition by the language designers that in some cases implementation flexibility provides a worthwhile benefit for language translators; this usage is not a defect in the language.

6.5.4 Applicable language characteristics

This vulnerability is intended to be applicable to languages with the following characteristics:

- Languages whose specification allows some variation in how a translator handles some construct, where reliance on one form of this variation can result in differences in external program behaviour.
- Language implementations may not be required to provide a mechanism for controlling implementation-defined behaviour.

6.5.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Document the set of implementation-defined features an application depends upon, so that upon a change of translator, development tools, or target configuration it can be ensured that those dependencies are still met.
- Ensure that a specific use of a construct having implementation-defined behaviour produces an external behaviour that is the same for all of the possible behaviours permitted by the language specification.
- Only use a language implementation whose implementation-defined behaviours are within a known subset of implementation-defined behaviours. The known subset should be chosen so that the 'same external behaviour' condition described above is met.

- 1 • Create highly visible documentation (perhaps at the start of a source file) that the default implementation-
- 2 defined behaviour is changed within the current file.
- 3 • When developing coding guidelines for a specific language all constructs that have implementation-defined
- 4 behaviour shall be documented and for each construct, the situations where the set of possible behaviours
- 5 can vary shall be enumerated.
- 6 • When applying this guideline on a project the functionality provided by and for changing its implementation-
- 7 defined behaviour shall be documented.
- 8 • Verify code behaviour using at least two different compilers with two different technologies.

9 **6.5.6 Implications for standardization**

10 In future standardization activities, the following items should be considered:

- 11 • Portability guidelines for a specific language should provide a list of common implementation-defined
- 12 behaviours.
- 13 • Language specifiers should enumerate all the cases of implementation-defined behaviour.
- 14 • Language designers should provide language directives that optionally disable obscure language features.

15 **6.6 Deprecated Language Features [MEM]**

16 **6.6.1 Description of application vulnerability**

17 All code should conform to the current standard for the respective language. In reality though, a language standard
 18 may change during the creation of a software system or suitable compilers and development environments may not
 19 be available for the new standard for some period of time after the standard is published. To smooth the process of
 20 evolution, features that are no longer needed or which serve as the root cause of or contributing factor for safety or
 21 security problems are often deprecated to temporarily allow their continued use but to indicate that those features
 22 may be removed in the future. The deprecation of a feature is a strong indication that it should not be used. Other
 23 features, although not formally deprecated, are rarely used and there exist other more common ways of expressing
 24 the same function. Use of these rarely used features can lead to problems when others are assigned the task of
 25 debugging or modifying the code containing those features.

26 **6.6.2 Cross reference**

27 JSF AV Rules: 8 and 11
 28 MISRA C 2004: 1.1, 4.2, and 20.10
 29 MISRA C++ 2008: 1-0-1, 2-3-1, 2-5-1, 2-7-1, 5-2-4, and 18-0-2
 30 Ada Quality and Style Guide: 7.1.1

31 **6.6.3 Mechanism of failure**

32 Most languages evolve over time. Sometimes new features are added making other features extraneous.
 33 Languages may have features that are frequently the basis for security or safety problems. The deprecation of
 34 these features indicates that there is a better way of accomplishing the desired functionality. However, there is
 35 always a time lag between the acknowledgement that a particular feature is the source of safety or security
 36 problems, the decision to remove or replace the feature and the generation of warnings or error messages by
 37 compilers that the feature shouldn't be used. Given that software systems can take many years to develop, it is
 38 possible and even likely that a language standard will change causing some of the features used to be suddenly
 39 deprecated. Modifying the software can be costly and time consuming to remove the deprecated features.
 40 However, if the schedule and resources permit, this would be prudent as future vulnerabilities may result from
 41 leaving the deprecated features in the code. Ultimately the deprecated features will likely need to be removed
 42 when the features are removed.

43 **6.6.4 Applicable language characteristics**

44 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 1 • All languages that have standards, though some only have defacto standards.
- 2 • All languages that evolve over time and as such could potentially have deprecated features at some point.

3 **6.6.5 Avoiding the vulnerability or mitigating its effects**

4 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 5 • Adhere to the latest published standard for which a suitable compiler and development environment is
6 available.
- 7 • Avoid the use of deprecated features of a language.
- 8 • Stay abreast of language discussions in language user groups and standards groups on the Internet.
9 Discussions and meeting notes will give an indication of problem prone features that should not be used or
10 should be used with caution.

11 **6.6.6 Implications for standardization**

12 In future standardization activities, the following items should be considered:

- 13 • Obscure language features for which there are commonly used alternatives should be considered for
14 removal from the language standard.
- 15 • Obscure language features that have routinely been found to be the root cause of safety or security
16 vulnerabilities, or that are routinely disallowed in software guidance documents should be considered for
17 removal from the language standard.
- 18 • Language designers should provide language mechanisms that optionally disable deprecated language
19 features.

20 **6.7 Pre-processor Directives [NMP]**

21 **6.7.1 Description of application vulnerability**

22 Pre-processor replacements happen before any source code syntax check, therefore there is no type checking –
23 this is especially important in function-like macro parameters.

24 If great care is not taken in the writing of macros, the expanded macro can have an unexpected meaning. In many
25 cases if explicit delimiters are not added around the macro text and around all macro arguments within the macro
26 text, unexpected expansion is the result.

27 Source code that relies heavily on complicated pre-processor directives may result in obscure and hard to maintain
28 code since the syntax they expect may be different from the expressions programmers regularly expect in a given
29 programming language.

30 **6.7.2 Cross reference**

31 Holtzmann-8
32 JSF SV Rules: 26, 27, 28, 29, 30, 31, and 32
33 MISRA C 2004: 19.6, 19.7, 19.8, and 19.9
34 MISRA C++ 2008: 16-0-3, 16-0-4, and 16-0-5
35 CERT C guidelines: PRE01-C, PRE02-C, PRE10-C, and PRE31-C

36 **6.7.3 Mechanism of failure**

37 Readability and maintainability may be greatly decreased if pre-processing directives are used instead of language
38 features.

39 While static analysis can identify many problems early; heavy use of the pre-processor can limit the effectiveness
40 of many static analysis tools, which typically work on the pre-processed source code.

1 In many cases where complicated macros are used, the program does not do what is intended. For example:

2 define a macro as follows,

```
#define CD(x, y) (x + y - 1) / y
```

3 whose purpose is to divide. Then suppose it is used as follows

```
a = CD (b & c, sizeof (int));
```

4 which expands into

```
a = (b & c + sizeof (int) - 1) / sizeof (int);
```

5 which most times will not do what is intended. Defining the macro as

```
#define CD(x, y) ((x) + (y) - 1) / (y)
```

6 will provide the desired result.

7 **6.7.4 Applicable language characteristics**

8 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 9 • Languages that have a lexical-level pre-processor.
- 10 • Languages that allow unintended groupings of arithmetic statements.
- 11 • Languages that allow cascading macros.
- 12 • Languages that allow duplication of side effects.
- 13 • Languages that allow macros that reference themselves.
- 14 • Languages that allow nested macro calls.
- 15 • Languages that allow complicated macros.

16 **6.7.5 Avoiding the vulnerability or mitigating its effects**

17 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 18 • Where it is possible to achieve the desired functionality without the use of pre-processor directives, this
- 19 should be done in preference to the use of pre-processor directives.
- 20

21 **6.7.6 Implications for standardization**

22 In future standardization activities, the following items should be considered:

- 23 • Standards should reduce or eliminate dependence on lexical-level pre-processors for essential functionality
- 24 (such as conditional compilation).
- 25 • Standards should consider providing capabilities to inline functions and procedure calls, to reduce the need
- 26 for pre-processor macros.

27 **6.8 Choice of Clear Names [NAI]**

28 **6.8.1 Description of application vulnerability**

29 Humans sometimes choose similar or identical names for objects, types, aggregates of types, subprograms and
 30 modules. They tend to use characteristics that are specific to the native language of the software developer to aid
 31 in this effort, such as use of mixed-casing, underscores and periods, or use of plural and singular forms to support
 32 the separation of items with similar names. Similarly, development conventions sometimes use casing for
 33 differentiation (for example, all uppercase for constants).

1 Human cognitive problems occur when different (but similar) objects, subprograms, types, or constants differ in
2 name so little that human reviewers are unlikely to distinguish between them, or when the system maps such
3 entities to a single entity.

4 Conventions such as the use of capitalization, and singular/plural distinctions may work in small and medium
5 projects, but there are a number of significant issues to be considered:

- 6 • Large projects often have mixed languages and such conventions are often language-specific.
- 7 • Many implementations support identifiers that contain international character sets and some language
8 character sets have different notions of casing and plurality.
- 9 • Different word-forms tend to be language and dialect specific, such as a pidgin, and may be meaningless
10 to humans that speak other dialects.

11
12 An important general issue is the choice of names that differ from each other negligibly (in human terms), for
13 example by differing by only underscores, (none, "_ " "_"), plurals ("s"), visually similar characters (such as "l" and
14 "1", "O" and "0"), or underscores/dashes ("-","_"). [There is also an issue where identifiers appear distinct to a
15 human but identical to the computer, such as FOO, Foo, and foo in some computer languages.] Character sets
16 extended with diacritical marks and non-Latin characters may offer additional problems. Some languages or their
17 implementations may pay attention to only the first n characters of an identifier.

18 The problems described above are different from overloading or overriding where the same name is used
19 intentionally (and documented) to access closely linked sets of subprograms. This is also different than using
20 reserved names which can lead to a conflict with the reserved use and the use of which may or may not be
21 detected at compile time.

22 Name confusion can lead to the application executing different code or accessing different objects than the writer
23 intended, or than the reviewers understood. This can lead to outright errors, or leave in place code that may
24 execute some time in the future with unacceptable consequences.

25 Although most such mistakes are unintentional, it is plausible that such usages can be intentional, if masking
26 surreptitious behaviour is a goal.

27 **6.8.2 Cross Reference**

28 JSF AV Rules: 48-56
29 MISRA C 2004: 1.4
30 CERT C guidelines: DCL02-C
31 Ada Quality and Style Guide: 3.2

32 **6.8.3 Mechanism of Failure**

33 Calls to the wrong subprogram or references to the wrong data element (that was missed by human review) can
34 result in unintended behaviour. Language processors will not make a mistake in name translation, but human
35 cognition limitations may cause humans to misunderstand, and therefore may be easily missed in human reviews.

36 **6.8.4 Applicable language characteristics**

37 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 38 • Languages with relatively flat name spaces will be more susceptible. Systems with modules, classes,
39 packages can use qualification to disambiguate names that originate from different parents.
- 40 • Languages that provide preconditions, postconditions, invariances and assertions or redundant coding of
41 subprogram signatures help to ensure that the subprograms in the module will behave as expected, but do
42 nothing if different subprograms are called.
- 43 • Languages that treat letter case as significant. Some languages do not differentiate between names with
44 differing case, while others do.

1 6.8.5 Avoiding the vulnerability or mitigating its effects

2 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 3 • Implementers can create coding standards that provide meaningful guidance on name selection and use.
4 Good language specific guidelines could eliminate most problems.
- 5 • Use static analysis tools to show the target of calls and accesses and to produce alphabetical lists of
6 names. Human review can then often spot the names that are sorted at an unexpected location or which
7 look almost identical to an adjacent name in the list.
- 8 • Use static tools (often the compiler) to detect declarations that are unused.
- 9 • Use languages with a requirement to declare names before use or use available tool or compiler options to
10 enforce such a requirement.

11 6.8.6 Implications for standardization

12 In future standardization activities, the following items should be considered:

- 13 • Languages that do not require declarations of names should consider providing an option that does impose
14 that requirement.

15 6.9 Choice of Filenames and other External Identifiers [AJN]

16 6.9.1 Description of application vulnerability

17 Interfacing with the directory structure or other external identifiers on a system on which software executes is very
18 common. Differences in the conventions used by operating systems can result in significant changes in behaviour
19 when the same program is executed under different operating systems. For instance, the directory structure,
20 permissible characters, case sensitivity, and so forth can vary among operating systems and even among
21 variations of the same operating system. For example, Microsoft XP prohibits “/?:&”*“<>|#%”; but UNIX, Linux, and
22 OS X operating systems allow any character except for the reserved character ‘/’ to be used in a filename.

23 Some operating systems are case sensitive while others are not. On non-case sensitive operating systems,
24 depending on the software being used, the same filename could be displayed, as “filename”, “Filename” or
25 “FILENAME” and all would refer to the same file.

26 Some operating systems, particularly older ones, only rely on the significance of the first n characters of the file
27 name. n can be unexpectedly small, such as the first 8 characters in the case of Win16 architectures which would
28 cause “filename1”, “filename2” and “filename3” to all map to the same file.

29 Variations in the filename, named resource or external identifier being referenced can be the basis for various kinds
30 of problems. Such mistakes or ambiguity can be unintentional, or intentional, and in either case they can be
31 potentially exploited, if surreptitious behaviour is a goal.

32 6.9.2 Cross Reference

33 JSF AV Rules: 46, 51, 53, 54, 55, and 56
34 MISRA C 2004: 1.4 and 5.1
35 CERT C guidelines: MSC09-C and MSC10-C

36 6.9.3 Mechanism of Failure

37 The wrong named resource, such as a file, may be used within a program in a form that provides access to a
38 resource that was not intended to be accessed. Attackers could exploit this situation to intentionally misdirect
39 access of a named resource to another named resource.

6.9.4 Applicable language characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Any language providing for use of an *API* (Application Programming Interface) for external access of resources with varied naming conventions. In practice, this means all languages.
- A particular language interface to a system should be consistent in its processing of filenames or external identifiers. Consistency is only the first consideration. Even though it is consistent, it may consistently do something that is unexpected by the developer of the software interfacing with the system.

6.9.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Where possible, use an API that provides a known common set of conventions for naming and accessing external resources, such as POSIX, ISO/IEC 9945:2003 (IEEE Std 1003.1-2001).
- Analyze the range of intended target systems, develop a suitable API for dealing with them, and document the analysis.
- Ensure that programs adapt their behaviour to the platform on which they are executing, so that only the intended resources are accessed. The means that information on such characteristics as the directory separator string and methods of accessing parent directories need to be parameterized and not exist as fixed strings within a program.
- Avoid creating resource names that are longer than the guaranteed unique length of all potential target platforms.
- Avoid creating resources, which are differentiated only by the case in their names.

6.9.6 Implications for standardization

In future standardization activities, the following items should be considered:

- Language APIs for interfacing with external identifiers should be compliant with ISO/IEC 9945:2003 (IEEE Std 1003.1-2001).

6.10 Unused Variable [XYR]

6.10.1 Description of application vulnerability

A variable's value is assigned but never used, making it a dead store. As a variant, a variable is declared but neither read nor written to in the program, making it an unused variable. This type of error suggests that the design has been incompletely or inaccurately implemented.

Unused variables by themselves are innocuous, but can be combined with other vulnerabilities such as index bounds errors and buffer overflows and may mask errors or provide hidden channels.

6.10.2 Cross reference

CWE:

563. Unused Variable

MISRA C++ 2008: 0-1-4 and 0-1-6

CERT C guidelines: MSC13-C

6.10.3 Mechanism of failure

A variable is declared, but never used. It is likely that the variable is simply vestigial, but it is also possible that the unused variable points out a bug. This is likely to suggest that the design has been incompletely or inaccurately implemented.

1 A variable is assigned a value but this value is never used thereafter. The assignment is then generally referred to
 2 as a dead store. Note that this may be acceptable if the variable is a volatile variable, for which the assignment of a
 3 value triggers some external event.

4 A dead store is indicative of careless programming or of a design or coding error; either the use of the value was
 5 forgotten (almost certainly an error) or the assignment was performed even though it was not needed (unless there
 6 is a justification for it).

7 An unused variable or a dead store is very unlikely to be the cause of a vulnerability. However, since compilers
 8 diagnose unused variables routinely and dead stores occasionally, their presence is often an indication that
 9 compiler warnings are either suppressed or are being ignored by programmers. This observation does not hold for
 10 automatically generated code, where it is commonplace to find unused variables and dead stores, introduced to
 11 keep the generation process simple and uniform.

12 **6.10.4 Applicable language characteristics**

13 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 14 • Dead stores are possible in any programming language that provides assignment. (Pure functional
 15 languages do not have this issue.)
- 16 • Unused variables (in the technical sense above) are possible only in languages that provide variable
 17 declarations.

18 **6.10.5 Avoiding the vulnerability or mitigating its effects**

19 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 20 • Enable detection of unused variables and dead stores in the compiler. The default setting may be to
 21 suppress these warnings.

22 **6.10.6 Implications for standardization**

23 In future standardization activities, the following items should be considered:

- 24 • Languages should consider requiring mandatory diagnostics for unused variables.

25 **6.11 Identifier Name Reuse [YOW]**

26 **6.11.1 Description of application vulnerability**

27 When distinct entities are defined in nested scopes using the same name it is possible that program logic will
 28 operate on an entity other than the one intended.

29 When it is not clear which identifier is used, the program could behave in ways that were not predicted by reading
 30 the source code. This can be found by testing, but circumstances can arise (such as the values of the same-named
 31 objects being mostly the same) where harmful consequences occur. This weakness can also lead to vulnerabilities
 32 such as hidden channels where humans believe that important objects are being rewritten or overwritten when in
 33 fact other objects are being manipulated.

34 For example, the innermost definition is deleted from the source, the program will continue to compile without a
 35 diagnostic being issued (but execution can produce unexpected results).

36 **6.11.2 Cross reference**

37 JSF AV Rules: 120 and 1359
 38 MISRA C 2004: 5.2, 5.5, 5.6, 5.7, 20.1, 20.2

1 MISRA C++ 2008: 2-10-2, 2-10-3, 2-10-4, 2-10-5, 2-10-6, 17-0-1, 17-0-2, and 17-0-3
 2 CERT C guidelines: DCL01-C and DCL32-C
 3 Ada Quality and Style Guide: 5.6.1 and 5.7.1

4 6.11.3 Mechanism of failure

5 Many languages support the concept of scope. One of the ideas behind the concept of scope is to provide a
 6 mechanism for the independent definition of identifiers that may share the same name.

7 For instance, in the following code fragment:

```
8
9     int some_var;
10
11     {
12         int t_var;
13         int some_var; /* definition in nested scope */
14
15         t_var=3;
16         some_var=2;
17     }
18
```

19 an identifier called `some_var` has been defined in different scopes.

20 If either the definition of `some_var` or `t_var` that occurs in the nested scope is deleted (for example, when the
 21 source is modified) it is necessary to delete all other references to the identifier's scope. If a developer deletes the
 22 definition of `t_var` but fails to delete the statement that references it, then most languages require a diagnostic to
 23 be issued (such as reference to undefined variable). However, if the nested definition of `some_var` is deleted but
 24 the reference to it in the nested scope is not deleted, then no diagnostic will be issued (because the reference
 25 resolves to the definition in the outer scope).

26 An example of how interpretations of a programming language can differ, in the following code fragment:

```
27     int j = 100;
28     {
29         for (int j = 0; j < 10; j++) ;
30         std::cout << j << std::endl; // What is the value of j
31     }
```

32 According to ISO 14882:2003 (C++) standard the value printed for `j` should be 100, but in some implementations
 33 that do not conform to the current version of the standard it will be 10, as the loop counter `j` remains in-scope after
 34 the end of the loop statement.

35 In some cases non-unique identifiers in the same scope can also be introduced through the use of identifiers
 36 whose common substring exceeds the length of characters the implementation considers to be distinct. For
 37 example, in the following code fragment:

```
38     extern int global_symbol_definition_lookup_table_a[100];
39     extern int global_symbol_definition_lookup_table_b[100];
```

40 the external identifiers are not unique on implementations where only the first 31 characters are significant. This
 41 situation only occurs in languages that allow multiple declarations of the same identifier (other languages require a
 42 diagnostic message to be issued). (See, Choice of Filenames and other External Identifiers [AJN].)

43 A related problem exists in languages that allow overloading or overriding of keywords or standard library function
 44 identifiers. Such overloading can lead to confusion about which entity is intended to be referenced.

1 Definitions for new identifiers should not use a name that is already visible within the scope containing the new
2 definition. Alternately, utilize language-specific facilities that check for and prevent inadvertent overloading of
3 names should be used.

4 **6.11.4 Applicable language characteristics**

5 This vulnerability is intended to be applicable to languages with the following characteristics:

- 6 • Languages that allow the same name to be used for identifiers defined in nested scopes.
- 7 • Languages where unique names can be transformed into non-unique names as part of the normal tool
8 chain.

9 **6.11.5 Avoiding the vulnerability or mitigating its effects**

10 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 11 • Ensure that a definition of an entity does not occur in a scope where a different entity with the same name
12 is accessible and can be used in the same context. A language-specific project coding convention can be
13 used to ensure that such errors are detectable with static analysis.
- 14 • Ensure that a definition of an entity does not occur in a scope where a different entity with the same name
15 is accessible and has a type that permits it to occur in at least one context where the first entity can occur.
- 16 • Use language features, if any, which explicitly mark definitions of entities that are intended to hide other
17 definitions.
- 18 • Develop or use tools that identify name collisions or reuse when truncated versions of names cause
19 conflicts.
- 20 • Ensure that all identifiers differ within the number of characters considered to be significant by the
21 implementations that are likely to be used, and document all assumptions.
22

23 **6.11.6 Implications for standardization**

24 In future standardization activities, the following items should be considered:

- 25 • Languages should require mandatory diagnostics for variables with the same name in nested scopes.
- 26 • Languages should require mandatory diagnostics for variable names that exceed the length that the
27 implementation considers unique.
- 28 • Languages should consider requiring mandatory diagnostics for overloading or overriding of keywords or
29 standard library function identifiers.

30 **6.12 Namespace Issues [BJL]**

31 **6.12.1 Description of Application Vulnerability**

32 If a language provides separate, non-hierarchical namespaces and a means to make names declared in these
33 name spaces directly visible to an application, the potential of unintentional and possible disastrous change in
34 application behavior can arise, when names are added to a namespace during maintenance.

35 Namespaces include constructs like packages, modules, libraries, classes or any other means of grouping
36 declarations for import into other program units.

37 **6.12.2 Cross References**

38 [None]

6.12.3 Mechanism of Failure

The failure is best illustrated by an example. Namespace N_1 provides the name A but not B ; Namespace N_2 provides the name B but not A . The application wishes to use A from N_1 and B from N_2 . At this point, there are no obvious issues. The application chooses (or needs to) import the namespaces to obtain names for direct usage, for an example.

Use N_1, N_2 ; – presumed to make all names in N_1 and N_2 directly visible

```
... X := A + B;
```

The semantics of the above example are intuitive and unambiguous.

Later, during maintenance, the name B is added to N_1 . The change to the namespace usually implies a recompilation of dependent units. At this point, two declarations of B are applicable for the use of B in the above example.

Some languages try to disambiguate the above situation by stating preference rules in case of such ambiguity among names provided by different name spaces. If, in the above example, N_1 is preferred over N_2 , the meaning of the use of B changes silently, presuming that no typing error arises. Consequently the semantics of the program change silently and assuredly unintentionally, since the implementer of N_1 can not assume that all users of N_1 would prefer to take any declaration of B from N_1 rather than its previous namespace.

It does not matter what the preference rules actually is, as long as the namespaces are mutable. The above example is easily extended by adding A to N_2 to show a symmetric error situation for a different precedence rule.

If a language supports overloading of subprograms, the notion of “same name” used in the above example is extended to mean not only the same name, but also the same signature of the subprogram. For vulnerabilities associated with overloading and overriding, see Identifier Name Reuse [YOW]. In the context of namespaces, however, adding signature matching to the name binding process, merely extends the described problem from simple names to full signatures, but does not alter the mechanism or quality of the described vulnerability. In particular, overloading does not introduce more ambiguity for binding to declarations in different name spaces. This vulnerability not only creates unintentional errors. It also can be exploited maliciously, if the source of the application and of the namespaces is known to the aggressor and one of the namespaces is mutable by the attacker.

6.12.4 Applicable Language Characteristics

The vulnerability is applicable to languages with the following characteristics:

- Languages that support non-hierarchical separate name-spaces, have means to import all names of a namespace “wholesale” for direct use, and have preference rules to chose among multiple imported direct homographs. All three conditions need to be satisfied for the vulnerability to arise.

6.12.5 Avoiding the Vulnerability or Mitigating its Effects

Software developers can avoid the vulnerability or mitigate its effects in the following ways:

- Avoiding “wholesale” import directives
- Using only selective “single name” import directives or using fully qualified names (in both cases, provided that the language offers the respective capabilities)

6.12.6 Implications for Standardization

- Languages should not have preference rules among mutable namespaces. Ambiguities should be illegal and avoidable by the user, for example, by using names qualified by their originating namespace.

1 6.13 Type System [IHN]

2 6.13.1 Description of application vulnerability

3 When data values are converted from one data type to another, even when done intentionally, unexpected results
4 can occur.

5 6.13.2 Cross reference

6 JSF AV Rule: 148 and 183

7 MISRA C 2004: 6.1, 6.2, 6.3, 10.1, and 10.5

8 MISRA C++ 2008: 3-9-2, 5-0-3 to 5-0-14

9 CERT C guidelines: DCL07-C, DCL11-C, DCL35-C, EXP05-C and EXP32-C

10 Ada Quality and Style Guide: 3.4

11 6.13.3 Mechanism of failure

12 The *type* of a data object informs the compiler how values should be represented and which operations may be
13 applied. The *type system* of a language is the set of rules used by the language to structure and organize its
14 collection of types. Any attempt to manipulate data objects with inappropriate operations is a *type error*. A program
15 is said to be *type safe* (or *type secure*) if it can be demonstrated that it has no type errors [20].

16 Every programming language has some sort of type system. A language is *statically typed* if the type of every
17 expression is known at compile time. The type system is said to be *strong* if it guarantees type safety and *weak* if it
18 does not. There are strongly typed languages that are not statically typed because they enforce type safety with run
19 time checks [20].

20 In practical terms, nearly every language falls short of being strongly typed (in an ideal sense) because of the
21 inclusion of mechanisms to bypass type safety in particular circumstances. For that reason and because every
22 language has a different type system, this description will focus on taking advantage of whatever features for type
23 safety may be available in the chosen language.

24 Sometimes it is appropriate for a data value to be converted from one type to another *compatible* one. For
25 example, consider the following program fragment, written in no specific language:

```
26     float a;  
27     integer i;  
28     a := a + i;
```

29 The variable "i" is of integer type. It must be converted to the float type before it can be added to the data value.
30 An implicit conversion, as shown, is called coercion. If, on the other hand, the conversion must be explicit, for
31 example, "a := a + float(i)", then the conversion is called a *cast*.

32 Type *equivalence* is the strictest form of type compatibility; two types are equivalent if they are compatible without
33 using coercion or casting. Type equivalence is usually characterized in terms of *name type equivalence*—two
34 variables have the same type if they are declared in the same declaration or declarations that use the same type
35 name—or *structure type equivalence*—two variables have the same type if they have identical structures. There
36 are variations of these approaches and most languages use different combinations of them [22]. Therefore, a
37 programmer skilled in one language may very well code inadvertent type errors when using a different language.

38 It is desirable for a program to be type safe because the application of operations to operands of an inappropriate
39 type may produce unexpected results. In addition, the presence of type errors can reduce the effectiveness of
40 static analysis for other problems. Searching for type errors is a valuable exercise because their presence often
41 reveals design errors as well as coding errors. Many languages check for type errors—some at compile-time,
42 others at run-time. Obviously, compile-time checking is more valuable because it can catch errors that are not
43 executed by a particular set of test cases.

1 Making the most use of the type system of a language is useful in two ways. First, data conversions always bear
 2 the risk of changing the value. For example, a conversion from integer to float risks the loss of significant digits
 3 while the inverse conversion risks the loss of any fractional value. Conversion of an integer value from a type with a
 4 longer representation to a type with a shorter representation risks the loss of significant digits. This can produce
 5 particularly puzzling results if the value is used to index an array. Conversion of a floating-point value from a type
 6 with a longer representation to a type with a shorter representation risks the loss of precision. This can be
 7 particularly severe in computations where the number of calculations increase as a power of the problem size. (It
 8 should be noted that similar surprises can occur when an application is retargeted to a machine with different
 9 representations of numeric values.)

10 Second, a programmer can use the type system to increase the probability of catching design errors or coding
 11 blunders. For example, the following Ada fragment declares two distinct floating-point types:

```
12     type Celsius is new Float;
13     type Fahrenheit is new Float;
```

14 The declaration makes it impossible to add a value of type Celsius to a value of type Fahrenheit without explicit
 15 conversion.

16 6.13.4 Applicable language characteristics

17 This vulnerability is intended to be applicable to languages with the following characteristics:

- 18 • Languages that support multiple types and allow conversions between types.

19 6.13.5 Avoiding the vulnerability or mitigating its effects

20 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 21 • Take advantage of any facility offered by the programming language to declare distinct types and use any
 22 mechanism provided by the language processor and related tools to check for or enforce type compatibility.
- 23 • Use available language and tools facilities to preclude or detect the occurrence of coercion. If it is not
 24 possible, use human review to assist in searching for coercions.
- 25 • Avoid casting data values except when there is no alternative. Document such occurrences so that the
 26 justification is made available to maintainers.
- 27 • Use the most restricted data type that suffices to accomplish the job. For example, use an enumeration
 28 type to select from a limited set of choices (such as, a switch statement or the discriminant of a union type)
 29 rather than a more general type, such as integer. This will make it possible for tooling to check if all
 30 possible choices have been covered.
- 31 • Treat every compiler, tool, or run-time diagnostic concerning type compatibility as a serious issue. Do not
 32 resolve the problem by modifying the code by inserting an explicit cast, without further analysis; instead
 33 examine the underlying design to determine if the type error is a symptom of a deeper problem.
- 34 • Never ignore instances of coercion; if the conversion is necessary, convert it to a cast and document the
 35 rationale for use by maintainers.
- 36 • Analyze the problem to be solved to learn the magnitudes and/or the precisions of the quantities needed as
 37 auxiliary variables, partial results and final results.

38 6.13.6 Implications for standardization

39 In future standardization activities, the following items should be considered:

- 40 • Language specifiers should standardize on a common, uniform terminology to describe their type systems
 41 so that programmers experienced in other languages can reliably learn the type system of a language that
 42 is new to them.
- 43 • Provide a mechanism for selecting data types with sufficient capability for the problem at hand.
- 44 • Provide a way for the computation to determine the limits of the data types actually selected.

- 1 • Language implementers should consider providing compiler switches or other tools to provide the highest
2 possible degree of checking for type errors.

3 **6.14 Bit Representations[STR]**

4 **6.14.1 Description of application vulnerability**

5 Interfacing with hardware, other systems and protocols often requires access to one or more bits in a single
6 computer word, or access to bit fields that may cross computer words for the machine in question. Mistakes can be
7 made as to what bits are to be accessed because of the “endianness” of the processor (see below) or because of
8 miscellaneous. Access to those specific bits may affect surrounding bits in ways that compromise their integrity.
9 This can result in the wrong information being read from hardware, incorrect data or commands being given, or
10 information being mangled, which can result in arbitrary effects on components attached to the system.

11 **6.14.2 Cross reference**

12 JSF AV Rules 147, 154 and 155

13 MISRA C 2004: 3.5, 6.4, 6.5, and 12.7

14 MISRA C++ 2008: 5-0-21, 5-2-4 to 5-2-9, and 9-5-1

15 CERT C guidelines: EXP38-C, INT00-C, INT07-C, INT12-C, INT13-C, and INT14-C

16 Ada Quality and Style Guide: 7.6.1 through 7.6.9, and 7.3.1

17 **6.14.3 Mechanism of failure**

18 Computer languages frequently provide a variety of sizes for integer variables. Languages may support short,
19 integer, long, and even big integers. Interfacing with protocols, device drivers, embedded systems, low level
20 graphics or other external constructs may require each bit or set of bits to have a particular meaning. Those bit
21 sets may or may not coincide with the sizes supported by a particular language implementation. When they do not,
22 it is common practice to pack all of the bits into one word. Masking and shifting of the word using powers of two to
23 pick out individual bits or using sums of powers of 2 to pick out subsets of bits (for example, using $28=2^2+2^3+2^4$ to
24 create the mask 11100 and then shifting 2 bits) provides a way of extracting those bits. Knowledge of the
25 underlying bit storage is usually not necessary to accomplish simple extractions such as these. Problems can arise
26 when programmers mix their techniques to reference the bits or output the bits. Problems can arise when
27 programmers mix arithmetic and logical operations to reference the bits or output the bits. The storage ordering of
28 the bits may not be what the programmer expects.

29 Packing of bits in an integer is not inherently problematic. However, an understanding of the intricacies of bit level
30 programming must be known. Some computers or other devices store the bits left to right while others store them
31 right to left. The type of storage can cause problems when interfacing with external devices that expect the bits in
32 the opposite order. One problem arises when assumptions are made when interfacing with external constructs and
33 the ordering of the bits or words are not the same as the receiving entity. Programmers may inadvertently use the
34 sign bit in a bit field and then may not be aware that an arithmetic shift (sign extension) is being performed when
35 right shifting causing the sign bit to be extended into other fields. Alternatively, a left shift can cause the sign bit to
36 be one. Bit manipulations can also be problematic when the manipulations are done on binary encoded records
37 that span multiple words. The storage and ordering of the bits must be considered when doing bitwise operations
38 across multiple words as bytes may be stored in big endian or little endian format.

39 **6.14.4 Applicable language characteristics**

40 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 41 • Languages that allow bit manipulations.

42 **6.14.5 Avoiding the vulnerability or mitigating its effects**

43 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 1 • Any assumption about bit ordering should be explicitly documented.
- 2 • The way bit ordering is done on the host system and on the systems with which the bit manipulations will
- 3 be interfaced should be understood.
- 4 • Bit fields should be used in languages that support them.
- 5 • Bit operators should not be used on signed operands.
- 6 • Localize and document the code associated with explicit manipulation of bits and bit fields.

7 **6.14.6 Implications for standardization**

- 8 • For languages that are commonly used for bit manipulations, an API for bit manipulations that is
- 9 independent of word size and machine instruction set should be defined and standardized.

10 **6.15 Floating-point Arithmetic [PLF]**

11 **6.15.1 Description of application vulnerability**

12 Most real numbers cannot be represented exactly in a computer. To represent real numbers, most computers use
 13 ANSI/IEEE Std 754 [28]. The bit representation for a floating-point number can vary from compiler to compiler and
 14 on different platforms. Relying on a particular representation can cause problems when a different compiler is used
 15 or the code is reused on another platform. Regardless of the representation, many real numbers can only be
 16 approximated since representing the real number using a binary representation would require an endlessly
 17 repeating string of bits or more binary digits than are available for representation. Therefore it should be assumed
 18 that a floating-point number is only an approximation, even though it may be an extremely good one. Floating-point
 19 representation of a real number or a conversion to floating-point can cause surprising results and unexpected
 20 consequences to those unaccustomed to the idiosyncrasies of floating-point arithmetic.

21 Algorithms that use floating point can have anomalous behaviour when used with certain values. The most
 22 common results are erroneous results or algorithms that never terminate for certain segments of the numeric
 23 domain, or for isolated values.

24 **6.15.2 Cross reference**

25 JSF AV Rules: 146, 147, 184, 197, and 202
 26 MISRA C 2004: 1.5, 12.12, 13.3, and 13.4
 27 MISRA C++ 2008: 0-4-3, 3-9-3, and 6-2-2
 28 CERT C guidelines: FLP00-C, FP01-C, FLP02-C and FLP30-C
 29 Ada Quality and Style Guide: 5.5.6 and 7.2.1 through 7.2.8

30 **6.15.3 Mechanism of failure**

31 Floating-point numbers are generally only an approximation of the actual value. In the base 10 world, the value of
 32 $1/3$ is 0.333333... The same type of situation occurs in the binary world, but numbers that can be represented with
 33 a limited number of digits in base 10, such as $1/10=0.1$ become endlessly repeating sequences in the binary world.
 34 So $1/10$ represented as a binary number is:

35 0.0001100110011001100110011001100110011001100110011001100110011...

36 Which is $0*1/2 + 0*1/4 + 0*1/8 + 1*1/16 + 1*1/32 + 0*1/64...$ and no matter how many digits are used, the
 37 representation will still only be an approximation of $1/10$. Therefore when adding $1/10$ ten times, the final result
 38 may or may not be exactly 1.

39 Accumulating floating point values through the repeated addition of values, particularly relatively small values, can
 40 provide unexpected results. Using an accumulated value to terminate a loop can result in an unexpected number of
 41 iterations. Rounding and truncation can cause tests of floating-point numbers against other values to yield
 42 unexpected results. Another cause of floating point errors is reliance upon comparisons of floating point values or
 43 the comparison of a floating point value with zero. Tests of equality/inequality can vary due to propagation or

1 conversion errors. Differences in magnitudes of floating-point numbers can result in no change of a very large
2 floating-point number when a relatively small number is added to or subtracted from it.

3 Manipulating bits in floating-point numbers is also very implementation dependent. Though IEEE 754 is a
4 commonly used representation for floating-point data types, it is not universally used or required by all computer
5 languages. Some languages predate IEEE 754 and make the support for the standard optional. One IEEE 754
6 representation uses a 24-bit mantissa (including the sign bit) and an 8-bit exponent, but the number of bits
7 allocated to the mantissa and exponent can vary when using other representations, as can the particular
8 representation used for the mantissa and exponent. Even within IEEE 754, various alternative representations are
9 permitted for the “extended precision” format (from 80- to 128-bit representations, with or without a hidden bit).
10 Typically special representations are specified for positive and negative zero and infinity. Relying on a particular bit
11 representation is inherently problematic, especially when a new compiler is introduced or the code is reused on
12 another platform. The uncertainties arising from floating-point can be divided into uncertainty about the actual bit
13 representation of a given value (such as, big-endian or little-endian) and the uncertainty arising from the rounding
14 of arithmetic operations (for example, the accumulation of errors when imprecise floating-point values are used as
15 loop indices).

16 **6.15.4 Applicable language characteristics**

17 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 18 • All languages with floating-point variables can be subject to rounding or truncation errors.

19 **6.15.5 Avoiding the vulnerability or mitigating its effects**

20 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 21 • Do not use a floating-point expression in a Boolean test for equality. Instead, use coding that determines
22 the difference between the two values to determine whether the difference is acceptably small enough so
23 that two values can be considered equal. Note that if the two values are very large, the “small enough”
24 difference can be a very large number.
- 25 • Use library functions with known numerical characteristics whenever possible.
- 26 • Unless the use of floating-point is simple, an expert in numerical analysis should check the stability and
27 accuracy of the algorithm employed.
- 28 • Avoid the use of a floating-point variable as a loop counter. If necessary to use a floating-point value as a
29 loop control, use inequality to determine the loop control (that is, <, <=, > or >=).
- 30 • Understand the floating-point format used to represent the floating-point numbers. This will provide some
31 understanding of the underlying idiosyncrasies of floating-point arithmetic.
- 32 • Manipulating the bit representation of a floating-point number should not be done except with built-in
33 language operators and functions that are designed to extract the mantissa and exponent.
- 34 • Do not use floating-point for exact values such as monetary amounts. Use floating-point only when
35 necessary such as for fundamentally inexact values such as measurements.
- 36 • Consider the use of decimal floating-point facilities when available.

37 **6.15.6 Implications for standardization**

38 In future standardization activities, the following items should be considered:

- 39 • Languages that do not already adhere to or only adhere to a subset of ANSI/IEEE 754 should consider
40 adhering completely to the standard. Examples of standardization that should be considered:
 - 41 ○ C, which predates ANSI/IEEE 754 and currently has it as optional in the 9899:1999 version,
42 should consider requiring ANSI/IEEE 754 for floating-point arithmetic.
 - 43 ○ Java should consider fully adhering to ANSI/IEEE 754 instead of a subset.
- 44 • Languages should consider providing a means to generate diagnostics for code that attempts to test
45 equality of two floating point values.
- 46 • Languages should consider standardizing their data type to ISO/IEC 10967-1:1994 and ISO/IEC 10967-
47 2:2001.

6.16 Enumerator Issues [CCB]

6.16.1 Description of application vulnerability

Enumerations are a finite list of named entities that contain a fixed mapping from a set of names to a set of integral values (called the representation) and an order between the members of the set. In some languages there are no other operations available except order, equality, first, last, previous, and next; in others the full underlying representation operators are available, such as integer “+” and “-” and bit-wise operations.

Most languages that provide enumeration types also provide mechanisms to set non-default representations. If these mechanisms do not enforce whole-type operations and check for conflicts then some members of the set may not be properly specified or may have the wrong mappings. If the value-setting mechanisms are positional only, then there is a risk that improper counts or changes in relative order will result in an incorrect mapping.

For arrays indexed by enumerations with non-default representations, there is a risk of structures with holes, and if those indexes can be manipulated numerically, there is a risk of out-of-bound accesses of these arrays.

Most of these errors can be readily detected by static analysis tools with appropriate coding standards, restrictions and annotations. Similarly mismatches in enumeration value specification can be detected statically. Without such rules, errors in the use of enumeration types are computationally hard to detect statically as well as being difficult to detect by human review.

6.16.2 Cross reference

JSF AV Rule: 145
 MISRA C 2004: 9.3, 9.2, and 9.3
 MISRA C++ 2008: 8-5-3
 CERT C guidelines: INT09-C
 Holtzmann rule 6.
 Ada Quality and Style Guide: 3.4.2

6.16.3 Mechanism of failure

As a program is developed and maintained the list of items in an enumeration often changes in three basic ways: new elements are added to the list; order between the members of the set often changes; and representation (the map of values of the items) change. Expressions that depend on the full set or specific relationships between elements of the set can create value errors that could result in wrong results or in unbounded behaviours if used as array indices.

Improperly mapped representations can result in some enumeration values being unreachable, or may create “holes” in the representation where undefinable values can be propagated.

If arrays are indexed by enumerations containing nondefault representations, some implementations may leave space for values that are unreachable using the enumeration, with a possibility of unnecessarily large memory allocations or a way to pass information undetected (hidden channel).

When enumerators are set and initialized explicitly and the language permits incomplete initializers, then changes to the order of enumerators or the addition or deletion of enumerators can result in the wrong values being assigned or default values being assigned improperly. Subsequent indexing or switch/case statements can result in illegal accesses and possibly unbounded behaviours.

6.16.4 Applicable language Characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Languages that do not require full coverage of an enumeration in a switch/case statement.
- Languages that provide a default choice in a switch/case statement.

- Languages that permit incomplete mappings between enumerator specification and value assignment, or that provide a positional-only mapping require additional static analysis tools and annotations to help identify the complete mapping of every literal to its value.
- Languages that provide a trivial mapping to a type such as integer require additional static analysis tools to prevent mixed type errors. They also cannot prevent illegal values from being placed into variables of such enumerator types. For example:

```
enum Directions {back, forward, stop};
enum Directions a = forward, b = stop, c = a+b;
```

In this example, `c` may have a value not defined by the enumeration, and any further use as that enumeration will lead to erroneous results.

- Some languages provide no enumeration capability, leaving it to the programmer to define named constants to represent the values and ranges.

6.16.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Use static analysis tools that will detect inappropriate use of enumerators, such as using them as integers or bit maps, and that detect enumeration definition expressions that are incomplete or incorrect. For languages with a complete enumeration abstraction this is the compiler.
- When a language requires full coverage of an enumeration in a switch/case statement, a default choice should not be provided. For languages that do not require full coverage, then a default choice with a possible error generation should be provided to ensure that there is full coverage.
- When named syntax is available for representation setting, coverage analysis can eliminate the order issues and the incomplete coverage issues as long as no default choice is given.

6.16.6 Implications for standardization

In future standardization activities, the following items should be considered:

- Languages that currently permit arithmetic and logical operations on enumeration types could provide a mechanism to ban such operations program-wide.
- Languages that provide automatic defaults or that do not enforce static matching between enumerator definitions and initialization expressions could provide a mechanism to enforce such matching.

6.17 Numeric Conversion Errors [FLC]

6.17.1 Description of application vulnerability

Certain contexts in various languages may require exact matches with respect to types [25]:

```
aVar := anExpression
value1 + value2
foo(arg1, arg2, arg3, ... , argN)
```

Type conversion seeks to follow these exact match rules while allowing programmers some flexibility in using values such as: structurally-equivalent types in a name-equivalent language, types whose value ranges may be distinct but intersect (for example, subranges), and distinct types with sensible/meaningful corresponding values (for example, integers and floats). Explicit conversions are called *type casts*. An implicit type conversion between compatible but not necessarily equivalent types is called *type coercion*.

Numeric conversions can lead to a loss of data, if the target representation is not capable of representing the original value. For example, converting from an integer type to a smaller integer type can result in truncation if the

1 original value cannot be represented in the smaller size and converting a floating point to an integer can result in a
2 loss of precision or an out-of-range value.

3 Type conversion errors can lead to erroneous data being generated, algorithms that fail to terminate, array bounds
4 errors, and arbitrary program execution.

5 **6.17.2 Cross reference**

6 CWE:

7 192. Integer Coercion Error

8 MISRA C 2004: 10.1-10.6, 11.3-11.5, and 12.9

9 MISRA C++ 2008: 2-13-3, 5-0-3, 5-0-4, 5-0-5, 5-0-6, 5-0-7, 5-0-8, 5-0-9, 5-0-10, 5-2-5, 5-2-9, and 5-3-2

10 CERT C guidelines: FLP34-C, INT02-C, INT08-C, INT31-C, and INT35-C

11 **6.17.3 Mechanism of failure**

12 Numeric conversion errors results in data integrity issues, but they may also result in a number of safety and
13 security vulnerabilities.

14 Vulnerabilities typically occur when appropriate range checking is not performed, and unanticipated values are
15 encountered. These can result in safety issues, for example, then the Ariane 5 launcher failure occurred due to an
16 improperly handled conversion error resulting in the processor being shutdown [23].

17 Conversion errors can also result in security issues. An attacker may input a particular numeric value to exploit a
18 flaw in the program logic. The resulting erroneous value may then be used as an array index, a loop iterator, a
19 length, a size, state data, or in some other security critical manner. For example, a truncated integer value may be
20 used to allocate memory, while the actual length is used to copy information to the newly allocated memory,
21 resulting in a buffer overflow [23].

22 Numeric type conversion errors often lead to undefined states of execution resulting in infinite loops or crashes. In
23 some cases, integer type conversion errors can lead to exploitable buffer overflow conditions, resulting in the
24 execution of arbitrary code. Integer type conversion errors result in an incorrect value being stored for the variable
25 in question.

26 **6.17.4 Applicable language characteristics**

27 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 28 • Languages that perform implicit type conversion (coercion).
- 29 • Weakly typed languages that do not strictly enforce type rules.
- 30 • Languages that support logical, arithmetic, or circular shifts on integer values.
- 31 • Languages that do not generate exceptions on problematic conversions.

32 **6.17.5 Avoiding the vulnerability or mitigating its effects**

33 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 34 • The first line of defense against integer vulnerabilities should be range checking, either explicitly or through
35 strong typing. All integer values originating from a source that is not trusted should be validated for
36 correctness. However, it is difficult to guarantee that multiple input variables cannot be manipulated to
37 cause an error to occur in some operation somewhere in a program [23].
- 38 • An alternative or ancillary approach is to protect each operation. However, because of the large number of
39 integer operations that are susceptible to these problems and the number of checks required to prevent or
40 detect exceptional conditions, this approach can be prohibitively labor intensive and expensive to
41 implement.
- 42 • A language that generates exceptions on erroneous data conversions might be chosen. Design objects
43 and program flow such that multiple or complex casts are unnecessary. Ensure that any data type casting
44 that you must use is entirely understood to reduce the plausibility of error in use.

- 1 • The use of static analysis can often identify whether or not unacceptable numeric conversions will occur.

2 Verifiably in-range operations are often preferable to treating out of range values as an error condition because the
3 handling of these errors has been repeatedly shown to cause denial-of-service problems in actual applications.
4 Faced with a numeric conversion error, the underlying computer system may do one of two things: (a) signal some
5 sort of error condition, or (b) produce a numeric value that is within the range of representable values on that
6 system. The latter semantics may be preferable in some situations in that it allows the computation to proceed, thus
7 avoiding a denial-of-service attack. However, it raises the question of what numeric result to return to the user.

8 A recent innovation from ISO/IEC TR 24731-1 [26] is the definition of the `rsize_t` type for the C programming
9 language. Extremely large object sizes are frequently a sign that an object's size was calculated incorrectly. For
10 example, negative numbers appear as very large positive numbers when converted to an unsigned type like
11 `size_t`. Also, some implementations do not support objects as large as the maximum value that can be
12 represented by type `size_t`. For these reasons, it is sometimes beneficial to restrict the range of object sizes to
13 detect programming errors. For implementations targeting machines with large address spaces, it is recommended
14 that `R_SIZE_MAX` be defined as the smaller of the size of the largest object supported or $(SIZE_MAX \gg 1)$, even
15 if this limit is smaller than the size of some legitimate, but very large, objects. Implementations targeting machines
16 with small address spaces may wish to define `R_SIZE_MAX` as `SIZE_MAX`, which means that there is no object size
17 that is considered a runtime-constraint violation.

18 6.17.6 Implications for standardization

19 In future standardization activities, the following items should be considered:

- 20 • Languages should consider providing means similar to the ISO/IEC TR 24731-1 definition of `rsize_t` type
21 for C to restrict object sizes so as to expose programming errors.
22 • Languages should consider making all type conversions explicit or at least generating warnings for implicit
23 conversions where loss of data might occur.

24 6.18 String Termination [CJM]

25 6.18.1 Description of application vulnerability

26 Some programming languages use a termination character to indicate the end of a string. Relying on the
27 occurrence of the string termination character without verification can lead to either exploitation or unexpected
28 behaviour.

29 6.18.2 Cross reference

30 CWE:
31 170. Improper Null Termination
32 CERT C guidelines: STR03-C, STR31-C, STR32-C, and STR36-C

33 6.18.3 Mechanism of failure

34 String termination errors occur when the termination character is solely relied upon to stop processing on the string
35 and the termination character is not present. Continued processing on the string can cause an error or potentially
36 be exploited as a buffer overflow. This may occur as a result of a programmer making an assumption that a string
37 that is passed as input or generated by a library contains a string termination character when it does not.

38 Programmers may forget to allocate space for the string termination character and expect to be able to store an `n`
39 length character string in an array that is `n` characters long. Doing so may work in some instances depending on
40 what is stored after the array in memory, but it may fail or be exploited at some point.

6.18.4 Applicable language characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Languages that use a termination character to indicate the end of a string.
- Languages that do not do bounds checking when accessing a string or array.

6.18.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Do not rely solely on the string termination character.
- Use library calls that do not rely on string termination characters such as `strncpy` instead of `strcpy` in the standard C library.

6.18.6 Implications for standardization

In future standardization activities, the following items should be considered:

- Eliminating library calls that make assumptions about string termination characters.
- Checking bounds when an array or string is accessed.
- Specifying a string construct that does not need a string termination character.

6.19 Boundary Beginning Violation [XYX]

6.19.1 Description of application vulnerability

A buffer underwrite condition occurs when an array is indexed outside its lower bounds, or pointer arithmetic results in an access to storage that occurs before the beginning of the intended object.

6.19.2 Cross reference

CWE:

124. Boundary Beginning Violation ('Buffer Underwrite')

129. Unchecked Array Indexing

JSF AV Rule: 25

MISRA C 2004: 21.1

MISRA C++ 2008: 5-0-15 to 5-0-18

CERT C guidelines: ARR30-C, ARR32-C, and ARR38-C

6.19.3 Mechanism of failure

There are several kinds of failures (in all cases an exception may be raised if the accessed location is outside of some permitted range):

- A read access will return a value that has no relationship to the intended value, such as, the value of another variable or uninitialized storage.
- An out-of-bounds read access may be used to obtain information that is intended to be confidential.
- A write access will not result in the intended value being updated and may result in the value of an unrelated object (that happens to exist at the given storage location) being modified.
- When an array has been allocated storage on the stack an out-of-bounds write access may modify internal runtime housekeeping information (for example, a function's return address) which might change a programs control flow.

1 **6.19.4 Applicable language characteristics**

2 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 3 • Languages that do not detect and prevent an array being accessed outside of its declared bounds (either
- 4 by means of an index or by pointer arithmetic).
- 5 • Languages that do not automatically allocate storage when accessing an array element for which storage
- 6 has not already been allocated.

7 **6.19.5 Avoiding the vulnerability or mitigating its effects**

8 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:.

- 9 • Use of implementation provided functionality to automatically check array element accesses and prevent
- 10 out-of-bounds accesses.
- 11 • Use of static analysis to verify that all array accesses are within the permitted bounds. Such analysis may
- 12 require that source code contain certain kinds of information, such as, that the bounds of all declared
- 13 arrays be explicitly specified, or that pre- and post-conditions be specified.
- 14 • Sanity checks should be performed on all calculated expressions used as an array index or for pointer
- 15 arithmetic.

16 Some guideline documents recommend only using variables having an unsigned type when indexing an array, on

17 the basis that an unsigned type can never be negative. This recommendation simply converts an indexing

18 underflow to an indexing overflow because the value of the variable will wrap to a large positive value rather than a

19 negative one. Also some language support arrays whose lower bound is greater than zero, so an index can be

20 positive and be less than the lower bound.

21 In the past the implementation of array bound checking has sometimes incurred what has been considered to be a

22 high runtime overhead (often because unnecessary checks were performed). It is now practical for translators to

23 perform sophisticated analysis that significantly reduces the runtime overhead (because runtime checks are only

24 made when it cannot be shown statically that no bound violations can occur).

25 **6.19.6 Implications for standardization**

26 In future standardization activities, the following items should be considered:

- 27 • Languages that use pointer types should consider specifying a standardized feature for a pointer type that
- 28 would enable array bounds checking, if such a pointer type is not already in the standard.

29 **6.20 Unchecked Array Indexing [XYZ]**

30 **6.20.1 Description of application vulnerability**

31 Unchecked array indexing occurs when a value is used as an index into an array without checking that it falls within

32 the acceptable index range.

33 **6.20.2 Cross reference**

34 CWE:

35 129. Unchecked Array Indexing

36 JSF AV Rules: 164 and 15

37 MISRA C 2004: 21.1

38 MISRA C++ 2008: 5-0-15 to 5-0-18

39 CERT C guidelines: ARR30-C, ARR32-C, ARR33-C, and ARR38-C

40 Ada Quality and Style Guide: 5.5.1, 5.5.2, 7.6.7, and 7.6.8

6.20.3 Mechanism of failure

A single fault could allow both an overflow and underflow of the array index. An index overflow exploit might use buffer overflow techniques, but this can often be exploited without having to provide "large inputs." Array index overflows can also trigger out-of-bounds read operations, or operations on the wrong objects; that is, "buffer overflows" are not always the result. Unchecked array indexing, depending on its instantiation, can be responsible for any number of related issues. Most prominent of these possible flaws is the buffer overflow condition. Due to this fact, consequences range from denial of service, and data corruption, to arbitrary code execution. The most common condition situation leading to unchecked array indexing is the use of loop index variables as buffer indexes. If the end condition for the loop is subject to a flaw, the index can grow or shrink unbounded, therefore causing a buffer overflow or underflow. Another common situation leading to this condition is the use of a function's return value, or the resulting value of a calculation directly as an index in to a buffer. Unchecked array indexing can result in the corruption of relevant memory and perhaps instructions, lead to the program halting, if the values are outside of the valid memory area. If the memory corrupted is data, rather than instructions, the system might continue to function with improper values. If the corrupted memory can be effectively controlled, it may be possible to execute arbitrary code, as with a standard buffer overflow.

Language implementations might or might not statically detect out of bound access and generate a compile-time diagnostic. At runtime the implementation might or might not detect the out-of-bounds access and provide a notification at runtime. The notification might be treatable by the program or it might not be. Accesses might violate the bounds of the entire array or violate the bounds of a particular index. It is possible that the former is checked and detected by the implementation while the latter is not. The information needed to detect the violation might or might not be available depending on the context of use. (For example, passing an array to a subroutine via a pointer might deprive the subroutine of information regarding the size of the array.)

Aside from bounds checking, some languages have ways of protecting against out-of-bounds accesses. Some languages automatically extend the bounds of an array to accommodate accesses that might otherwise have been beyond the bounds. However, this may or may not match the programmer's intent and can mask errors. Some languages provide for whole array operations that may obviate the need to access individual elements thus preventing unchecked array accesses.

6.20.4 Applicable language characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Languages that do not automatically bounds check array accesses.
- Languages that do not automatically extend the bounds of an array to accommodate array accesses.

6.20.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Include sanity checks to ensure the validity of any values used as index variables.
- The choice could be made to use a language that is not susceptible to these issues.
- When available, use whole array operations whenever possible.

6.20.6 Implications for standardization

In future standardization activities, the following items should be considered:

- Languages should consider providing compiler switches or other tools to check the size and bounds of arrays and their extents that are statically determinable.
- Languages should consider providing whole array operations that may obviate the need to access individual elements.
- Languages should consider the capability to generate exceptions or automatically extend the bounds of an array to accommodate accesses that might otherwise have been beyond the bounds.

1 **6.21 Unchecked Array Copying [XYW]**

2 **6.21.1 Description of application vulnerability**

3 A buffer overflow occurs when some number of bytes (or other units of storage) is copied from one buffer to
4 another and the amount being copied is greater than is allocated for the destination buffer.

5 **6.21.2 Cross reference**

6 CWE:

7 121. Stack-based Buffer Overflow

8 JSF AV Rule: 15

9 MISRA C 2004: 21.1

10 MISRA C++ 2008: 5-0-15 to 5-0-18

11 CERT C guidelines: ARR33-C and STR31-C

12 Ada Quaility and Style Guide: 7.6.7 and 7.6.8

13 **6.21.3 Mechanism of failure**

14 Many languages and some third party libraries provide functions that efficiently copy the contents of one area of
15 storage to another area of storage. Most of these libraries do not perform any checks to ensure that the copied
16 from/to storage area is large enough to accommodate the amount of data being copied.

17 The arguments to these library functions include the addresses of the contents of the two storage areas and the
18 number of bytes (or some other measure) to copy. Passing the appropriate combination of incorrect start
19 addresses or number of bytes to copy makes it possible to read or write outside of the storage allocated to the
20 source/destination area. When passed incorrect parameters the library function performs one or more unchecked
21 array index accesses, as described in Unchecked Array Indexing [XYZ].

22 **6.21.4 Applicable language characteristics**

23 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 24 • Languages that contain standard library functions for performing bulk copying of storage areas.
- 25 • The same range of languages having the characteristics listed in Unchecked Array Indexing [XYZ].

26 **6.21.5 Avoiding the vulnerability or mitigating its effects**

27 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 28 • Only use library functions that perform checks on the arguments to ensure no buffer overrun can occur
29 (perhaps by writing a wrapper for the Standard provided functions). Perform checks on the argument
30 expressions prior to calling the Standard library function to ensure that no buffer overrun will occur.
- 31 • Use static analysis to verify that the appropriate library functions are only called with arguments that do not
32 result in a buffer overrun. Such analysis may require that source code contain certain kinds of information,
33 for example, that the bounds of all declared arrays be explicitly specified, or that pre- and post-conditions
34 be specified as annotations or language constructs.

35 **6.21.6 Implications for standardization**

36 In future standardization activities, the following items should be considered:

- 37 • Languages should consider only providing libraries that perform checks on the parameters to ensure that
38 no buffer overrun can occur.
- 39 • Languages should consider providing full array assignment.

6.22 Buffer Overflow [XZB]

6.22.1 Description of application vulnerability

A buffer overflow arises when, due to unchecked array indexing or unchecked array copying, storage outside the buffer is accessed. Usually overflows describe the situation where such storage is then written. Depending on where the buffer is located, logically unrelated portions of the stack or the heap could be modified maliciously or unintentionally. Usually, buffer overflows describe accesses to contiguous memory beyond the end of the buffer data, as may arise when arrays are copied without length checks. However, accessing before the beginning of the buffer data is equally possible, dangerous and maliciously exploitable.

6.22.2 Cross reference

CWE:

122. Heap-based Buffer Overflow

JSF AV Rule: 15

MISRA C 2004: 21.1

MISRA C++ 2008: 5-0-15 to 5-0-18

CERT C guidelines: ARR33-C, STR31-C and MEM35-C

6.22.3 Mechanism of failure

Overwriting adjacent data (or data at arbitrarily computed locations) outside the area allocated for an array leads to value failures of the application. The program statements causing the buffer overflow are often difficult to find. But not only data storage can be corrupted. Buffer overflow may also inadvertently or even maliciously overwrite function pointers that may be in memory, pointing them to the attacker's code. Even in applications that do not explicitly use function pointers, the run-time will usually store function pointers in memory. For example, object methods in object-oriented languages are generally implemented using function pointers in data structures that are kept in memory. Since the consequence of a buffer overflow can be targeted to cause arbitrary code execution, this vulnerability can be used to subvert any security service.

6.22.4 Applicable language characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Copying of arrays can be done without an automatic length check ensuring that source and target locations are of the same size.
- Indexing of array elements can be done without an automatic check that the indexing is within the bounds of the array.
- Accesses might violate the physical bounds of the entire array or violate the logical bounds of a particular extent. The vulnerability is somewhat mitigated, if the former violation is checked for and detected by the implementation although the latter is not.
- Languages that provide bounds checking but permit the check to be suppressed.
- The bounds of an array are not automatically extended to accommodate accesses that might otherwise have been beyond the bounds. (This may or may not match the programmer's intent.)

6.22.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Use a language or compiler that performs automatic bounds checking on elements accesses and automatic length checking on copying entire arrays.
- Use an abstraction library to add checks on top of library functions that copy arrays without length checks.
- Checks that prevent overflows can be disabled in some languages to increase performance. This option should be used rarely.

- Implementation-defined checks that prevent overflows can be enabled in some languages that do not require such checks. This option should be used whenever feasible.

6.22.6 Implications for standardization

In future standardization activities, the following items should be considered:

- Languages should provide safe copying of arrays as built-in operation.
- Languages should consider only providing array copy routines in libraries that perform checks on the parameters to ensure that no buffer overrun can occur.
- Languages should perform automatic bounds checking on accesses to array elements. This capability may need to be disabled at times for performance reasons.

6.23 Pointer Casting and Pointer Type Changes [HFC]

6.23.1 Description of application vulnerability

The code produced for access via a data or function pointer requires that the type of the pointer is appropriate for the data or function being accessed. Otherwise undefined behaviour can occur. Specifically, “access via a data pointer” is defined to be “fetch or store indirectly through that pointer” and “access via a function pointer” is defined to be “invocation indirectly through that pointer.” The detailed requirements for what is meant by the “appropriate” type may vary among languages.

Even if the type of the pointer is appropriate for the access, erroneous pointer operations can still cause a fault.

6.23.2 Cross reference

CWE

136. Type Errors

188. Reliance on Data/Memory Layout

JSF AV Rules: 182 and 183

MISRA C 2004: 11.1, 11.2, 11.3, 11.4, and 11.5

MISRA C++ 2008: 5-2-2 to 5-2-9

CERT C guidelines: INT11-C and EXP36-A

Hatton 13: Pointer casts

Ada Quaility and Style Guide: 7.6.7 and 7.6.8

6.23.3 Mechanism of failure

If a pointer's type is not appropriate for the data or function being accessed, data can be corrupted or privacy can be broken by inappropriate read or write operation using the indirection provided by the pointer value. With a suitable type definition, large portions of memory can be maliciously or accidentally modified or read. Such modification of data objects will generally lead to value faults of the application. Modification of code elements such as function pointers or internal data structures for the support of object-orientation can affect control flow. This can make the code susceptible to targeted attacks by causing invocation via a pointer-to-function that has been manipulated to point to an attacker's payload.

6.23.4 Applicable language characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Pointers (and/or references) can be converted to different pointer types.
- Pointers to functions can be converted to pointers to data.

1 **6.23.5 Avoiding the vulnerability or mitigating its effects**

2 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 3 • Treat the compiler's pointer-conversion warnings as serious errors.
- 4 • Adopt programming guidelines (preferably augmented by static analysis) that restrict pointer conversions.
5 For example, consider the rules itemized above from JSF AV [5], CERT C [11], Hatton [18], or MISRA C
6 [12].
- 7 • Other means of assurance might include proofs of correctness, analysis with tools, verification techniques,
8 or other methods.

9 **6.23.6 Implications for standardization**

10 In future standardization activities, the following items should be considered:

- 11 • Languages should consider creating a mode that provides a runtime check of the validity of all accessed
12 objects before the object is read, written or executed.

13 **6.24 Pointer Arithmetic [RVG]**

14 **6.24.1 Description of application vulnerability**

15 Using pointer arithmetic incorrectly can result in addressing arbitrary locations, which in turn can cause a program
16 to behave in unexpected ways.

17 **6.24.2 Cross reference**

18 JSF AV Rule: 215
19 MISRA C 2004: 17.1, 17.2, 17.3, and 17.4
20 MISRA C++ 2008: 5-0-15 to 5-0-18
21 CERT C guidelines: EXP08-C

22 **6.24.3 Mechanism of failure**

23 Pointer arithmetic used incorrectly can produce:

- 24 • Addressing arbitrary memory locations, including buffer underflow and overflow.
- 25 • Arbitrary code execution.
- 26 • Addressing memory outside the range of the program.

27 **6.24.4 Applicable language characteristics**

28 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 29 • Languages that allow pointer arithmetic.

30 **6.24.5 Avoiding the vulnerability or mitigating its effects**

31 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 32 • Use pointer arithmetic only for indexing objects defined as arrays.
- 33 • Prefer indexing for accessing array elements rather than using pointer arithmetic.
- 34 • Limit pointer arithmetic calculations to the addition and subtraction of integers.

1 6.24.6 Implications for standardization

2 [None]

3 6.25 Null Pointer Dereference [XYH]

4 6.25.1 Description of application vulnerability

5 A null-pointer dereference takes place when a pointer with a value of `NULL` is used as though it pointed to a valid
6 memory location. This is a special case of accessing storage via an invalid pointer.

7 6.25.2 Cross reference

8 CWE:

9 476. NULL Pointer Dereference

10 JSF AV Rule 174

11 CERT C guidelines: EXP34-C

12 Ada Quality and Style Guide: 5.4.5

13 6.25.3 Mechanism of failure

14 When a pointer with a value of `NULL` is used as though it pointed to a valid memory location, then a null-pointer
15 dereference is said to take place. This can result in a segmentation fault, unhandled exception, or accessing
16 unanticipated memory locations.

17 6.25.4 Applicable language characteristics

18 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 19 • Languages that permit the use of pointers and that do not check the validity of the location being accessed
20 prior to the access.
- 21 • Languages that allow the use of a `NULL` pointer.

22 6.25.5 Avoiding the vulnerability or mitigating its effects

23 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 24 • Before dereferencing a pointer, ensure it is not equal to `NULL`.

25 6.25.6 Implications for standardization

26 In future standardization activities, the following items should be considered:

- 27 • A language feature what would have pointers be checked for a `NULL` value before performing the access
28 should be considered.

29 6.26 Dangling Reference to Heap [XYK]

30 6.26.1 Description of application vulnerability

31 A dangling reference is a reference to an object whose lifetime has ended due to explicit deallocation or the stack
32 frame in which the object resided has been freed due to exiting the dynamic scope. The memory for the object may
33 be reused; therefore, any access through the dangling reference may affect an apparently arbitrary location of
34 memory, corrupting data or code.

1 This description concerns the former case, dangling references to the heap. The description of dangling references
2 to stack frames is [DCM]. In many languages references are called pointers; the issues are identical.

3 A notable special case of using a dangling reference is calling a deallocator, for example, `free()`, twice on the
4 same pointer value. Such a “Double Free” may corrupt internal data structures of the heap administration, leading
5 to faulty application behaviour (such as infinite loops within the allocator, returning the same memory repeatedly as
6 the result of distinct subsequent allocations, or deallocating memory legitimately allocated to another request since
7 the first `free()` call, to name but a few), or it may have no adverse effects at all.

8 Memory corruption through the use of a dangling reference is among the most difficult of errors to locate.

9 With sufficient knowledge about the heap management scheme (often provided by the OS (Operating System) or
10 run-time system), use of dangling references is an exploitable vulnerability, since the dangling reference provides a
11 method with which to read and modify valid data in the designated memory locations after freed memory has been
12 re-allocated by subsequent allocations.

13 6.26.2 Cross reference

14 CWE:

15 415. Double Free (Note that Double Free (415) is a special case of Use After Free (416))

16 416. Use After Free

17 MISRA C 2004: 17.1-6

18 MISRA C++ 2008: 0-3-1, 7-5-1, 7-5-2, 7-5-3, and 18-4-1

19 CERT C guidelines: MEM01-C, MEM30-C, and MEM31.C

20 Ada Quality and Style Guide: 5.4.5, 7.3.3, and 7.6.6

21 6.26.3 Mechanism of failure

22 The lifetime of an object is the portion of program execution during which storage is guaranteed to be reserved for
23 it. An object exists and retains its last-stored value throughout its lifetime. If an object is referred to outside of its
24 lifetime, the behaviour is undefined. Explicit deallocation of heap-allocated storage ends the lifetime of the object
25 residing at this memory location (as does leaving the dynamic scope of a declared variable). The value of a pointer
26 becomes indeterminate when the object it points to reaches the end of its lifetime. Such pointers are called
27 dangling references.

28 The use of dangling references to previously freed memory can have any number of adverse consequences —
29 ranging from the corruption of valid data to the execution of arbitrary code, depending on the instantiation and
30 timing of the deallocation causing all remaining copies of the reference to become dangling, of the system's reuse
31 of the freed memory, and of the subsequent usage of a dangling reference.

32 Like memory leaks and errors due to double de-allocation, the use of dangling references has two common and
33 sometimes overlapping causes:

- 34 • An error condition or other exceptional circumstances.
- 35 • Developer confusion over which part of the program is responsible for freeing the memory.

36 If a pointer to previously freed memory is used, it is possible that the referenced memory has been reallocated.
37 Therefore, assignment using the original pointer has the effect of changing the value of an unrelated variable. This
38 induces unexpected behaviour in the affected program. If the newly allocated data happens to hold a class
39 description, in an object-oriented language for example, various function pointers may be scattered within the heap
40 data. If one of these function pointers is overwritten with an address of malicious code, execution of arbitrary code
41 can be achieved.

42 6.26.4 Applicable language characteristics

43 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 1 • Languages that permit the use of pointers and that permit explicit deallocation by the developer or provide
- 2 for alternative means to reallocate memory still pointed to by some pointer value.
- 3 • Languages that permit definitions of constructs that can be parameterized without enforcing the
- 4 consistency of the use of parameter at compile time.

5 **6.26.5 Avoiding the vulnerability or mitigating its effects**

6 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 7 • Use an implementation that checks whether a pointer is used that designates a memory location that has
- 8 already been freed.
- 9 • Use a coding style that does not permit deallocation.
- 10 • In complicated error conditions, be sure that clean-up routines respect the state of allocation properly. If the
- 11 language is object-oriented, ensure that object destructors delete each chunk of memory only once.
- 12 Ensuring that all pointers are set to NULL once the memory they point to have been freed can be an
- 13 effective strategy. The utilization of multiple or complex data structures may lower the usefulness of this
- 14 strategy.
- 15 • Use a static analysis tool that is capable of detecting some situations when a pointer is used after the
- 16 storage it refers to is no longer a pointer to valid memory location.
- 17 • Allocating and freeing memory in different modules and levels of abstraction burdens the programmer with
- 18 tracking the lifetime of that block of memory. This may cause confusion regarding when and if a block of
- 19 memory has been allocated or freed, leading to programming defects such as double-free vulnerabilities,
- 20 accessing freed memory, or dereferencing NULL pointers or pointers that are not initialized. To avoid
- 21 these situations, it is recommended that memory be allocated and freed at the same level of abstraction,
- 22 and ideally in the same code module.

23 **6.26.6 Implications for standardization**

24 In future standardization activities, the following items should be considered:

- 25 • Implementations of the free function could tolerate multiple frees on the same reference/pointer or frees of
- 26 memory that was never allocated.
- 27 • Language specifiers should design generics in such a way that any attempt to instantiate a generic with
- 28 constructs that do not provide the required capabilities results in a compile-time error.
- 29 • For properties that cannot be checked at compile time, language specifiers should provide an assertion
- 30 mechanism for checking properties at run-time. It should be possible to inhibit assertion checking if
- 31 efficiency is a concern.
- 32 • A storage allocation interface should be provided that will allow the called function to set the pointer used to
- 33 NULL after the referenced storage is deallocated.

34 **6.27 Templates and Generics [SYM]**

35 **6.27.1 Description of application vulnerability**

36 Many languages provide a mechanism that allows objects and/or functions to be defined parameterized by type
 37 and then instantiated for specific types. In C++ and related languages, these are referred to as “templates”, and in
 38 Ada and Java, “generics”. To avoid having to keep writing ‘templates/generics’, in this section these will simply be
 39 referred to collectively as generics.

40 Used well, generics can make code clearer, more predictable and easier to maintain. Used badly, they can have
 41 the reverse effect, making code difficult to review and maintain, leading to the possibility of program error.

42 **6.27.2 Cross reference**

43 JSF AV Rules: 101, 102, 103, 104, and 105
 44 MISRA C 2004: 14-7-2, 14-8-1, and 14-8-2

1 MISRA C++ 2008: 14-6-1, 14-6-2, 14-7-1 to 14-7-3, 14-8-1, and 14-8-2
2 Ada Quality and Style Guide: 8.3.1 through 8.3.8, and 8.4.2

4 6.27.3 Mechanism of failure

5 The value of generics comes from having a single piece of code that supports some behaviour in a type
6 independent manner. This simplifies development and maintenance of the code. It should also assist in the
7 understanding of the code during review and maintenance, by providing the same behaviour for all types with which
8 it is instantiated.

9 Problems arise when the use of a generic actually makes the code harder to understand during review and
10 maintenance, by not providing consistent behaviour.

11 In most cases, the generic definition will have to make assumptions about the types it can legally be instantiated
12 with. For example, a sort function requires that the elements to be sorted can be copied and compared. If these
13 assumptions are not met, the result is likely to be a compiler error. For example if the sort function is instantiated
14 with a user defined type that doesn't have a relational operator. Where 'misuse' of a generic leads to a compiler
15 error, this can be regarded as a development issue, and not a software vulnerability.

16 Confusion, and hence potential vulnerability, can arise where the instantiated code is apparently illegal, but doesn't
17 result in a compiler error. For example, a generic class defines a set of members, a subset of which rely on a
18 particular property of the instantiation type (such as a generic container class with a sort member function, only the
19 sort function relies on the instantiating type having a defined relational operator). In some languages, such as C++,
20 if the generic is instantiated with a type that doesn't meet all the requirements but the program never subsequently
21 makes use of the subset of members that rely on the property of the instantiating type, the code will compile and
22 execute (for example, the generic container is instantiated with a user defined class that doesn't define a relational
23 operator, but the program never calls the sort member of this instantiation). When the code is reviewed the generic
24 class will appear to reference a member of the instantiating type that doesn't exist.

25 The problem as described in the two prior paragraphs can be reduced by a language feature (such as the *concepts*
26 language feature being designed by the C++ committee).

27 Similar confusion can arise if the language permits specific elements of a generic to be explicitly defined, rather
28 than using the common code, so that behaviour is not consistent for all instantiations. For example, for the same
29 generic container class, the sort member normally sorts the elements of the container into ascending order. In
30 languages such as C++, a 'special case' can be created for the instantiation of the generic with a particular type.
31 For example, the sort member for a 'float' container may be explicitly defined to provide different behaviour, say
32 sorting the elements into descending order. Specialization that doesn't affect the apparent behaviour of the
33 instantiation is not an issue. Again, for C++, there are some irregularities in the semantics of arrays and pointers
34 that can lead to the generic having different behaviour for different, but apparently very similar, types. In such
35 cases, specialization can be used to enforce consistent behaviour.

36 6.27.4 Applicable language characteristics

37 This vulnerability is intended to be applicable to languages with the following characteristics:

- 38 • Languages that permit definitions of objects or functions to be parameterized by type, for later instantiation
39 with specific types, such as:
 - 40 ○ Templates in C++
 - 41 ○ Generics in Ada, Java

42 6.27.5 Avoiding the vulnerability or mitigating its effects

43 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 44 • Document the properties of an instantiating type necessary for the generic to be valid.

- 1 • If an instantiating type has the required properties, the whole of the generic should be ensured to be valid,
2 whether actually used in the program or not.
- 3 • Preferably avoid, but at least carefully document, any 'special cases' where the generic instantiated with a
4 specific type doesn't behave as it does for other types.

5 **6.27.6 Implications for standardization**

6 In future standardization activities, the following items should be considered:

- 7 • Language specifiers should standardize on a common, uniform terminology to describe generics/templates
8 so that programmers experienced in one language can reliably learn and refer to the type system of
9 another language that has the same concept, but with a different name.
- 10 • Language specifiers should design generics in such a way that any attempt to instantiate a generic with
11 constructs that do not provide the required capabilities results in a compile-time error.
- 12 • Language specifiers should provide an assertion mechanism for checking properties at run-time, for those
13 properties that cannot be checked at compile time.. It should be possible to inhibit assertion checking if
14 efficiency is a concern.

15 **6.28 Inheritance [RIP]**

16 **6.28.1 Description of application vulnerability**

17 Inheritance, the ability to create enhanced and/or restricted object classes based on existing object classes can
18 introduce a number of vulnerabilities, both inadvertent and malicious. Because Inheritance allows the overriding of
19 methods of the parent class and because object oriented systems are designed to separate and encapsulate code
20 and data, it can be difficult to determine where in the hierarchy an invoked method is actually defined. Also, since
21 an overriding method does not need to call the method in the parent class that has been overridden, essential
22 initialization and manipulation of class data may be bypassed. This can be especially dangerous during constructor
23 and destructor methods.

24 Languages that allow multiple inheritance add additional complexities to the resolution of method invocations.
25 Different object brokerage systems may resolve the method identity to different classes, based on how the
26 inheritance tree is traversed.

27 **6.28.2 Cross reference**

28 JSF AV Rules: 86 to 97
29 MISRA C++ 2008: 0-1-12, 8-3-1, 10-1-1 to 10-1-3, and 10-3-1 to 10-3-3
30 Ada Quaility and Style Guide: 9 (complete section)

31 **6.28.3 Mechanism of failure**

32 The use of inheritance can lead to an exploitable application vulnerability or negatively impact system safety in
33 several ways:

- 34 • Execution of malicious redefinitions, this can occur through the insertion of a class into the class hierarchy
35 that overrides commonly called methods in the parent classes.
- 36 • Accidental redefinition, where a method is defined that inadvertently overrides a method that has already
37 been defined in a parent class.
- 38 • Accidental failure of redefinition, when a method is incorrectly named or the parameters are not defined
39 properly, and thus does not override a method in a parent class.
- 40 • Breaking of class invariants, this can be caused by redefining methods that initialize or validate class data
41 without including that initialization or validation in the overriding methods.

42 These vulnerabilities can increase dramatically as the complexity of the hierarchy increases, especially in the use
43 of multiple inheritance.

1 **6.28.4 Applicable language characteristics**

2 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 3
 - Languages that allow single and multiple inheritances.

4 **6.28.5 Avoiding the vulnerability or mitigating its effects**

5 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 6
 - Avoid the use of multiple inheritance whenever possible.
 - 7 • Provide complete documentation of all encapsulated data, and how each method affects that data for each object in the hierarchy.
 - 8 • Inherit only from trusted sources, and, whenever possible, check the version of the parent classes during compilation and/or initialization.
 - 9 • Provide a method that provides versioning information for each class.

12 **6.28.6 Implications for standardization**

13 In future standardization activities, the following items should be considered:

- 14
 - Language specification should include the definition of a common versioning method.
 - 15 • Compilers should provide an option to report the class in which a resolved method resides.
 - 16 • Runtime environments should provide a trace of all runtime method resolutions.

17 **6.29 Initialization of Variables [LAV]**

18 **6.29.1 Description of application vulnerability**

19 Reading a variable that has not been assigned a value appropriate to its type can cause unpredictable execution in the block that uses the value of the variable, and has the potential to export bad values to callers, or cause out-of-bounds memory accesses.

22 Uninitialized variable usage is frequently not detected until after testing and often when the code in question is delivered and in use, because happenstance will provide variables with adequate values (such as default data settings or accidental left-over values) until some other change exposes the defect.

25 Variables that are declared during module construction (by a class constructor, instantiation, or elaboration) may have alternate paths that can read values before they are set. This can happen in straight sequential code but is more prevalent when concurrency or co-routines are present, with the same impacts described above.

28 Another vulnerability occurs when compound objects are initialized incompletely, as can happen when objects are incrementally built, or fields are added under maintenance.

30 When possible and supported by the language, whole-structure initialization is preferable to field-by-field initialization statements, and named association is preferable to positional, as it facilitates human review and is less susceptible to failures under maintenance. For classes, the declaration and initialization may occur in separate modules. In such cases it must be possible to show that every field that needs an initial value receives that value, and to document ones that do not require initial values.

35 **6.29.2 Cross reference**

36 CWE:

37 457. Use of Uninitialized Variable

38 JSF AV Rules: 71, 143, and 147

39 MISRA C 2004: 9.1, 9.2, and 9.3

40 CERT C guidelines: DCL14-C and EXP33-C

1 MISRA C++ 2008: 8-5-1
2 Ada Quality and Style Guide: 5.9.6

3 **6.29.3 Mechanism of failure**

4 Uninitialized objects may have illegal values, legal but wrong values, or legal and dangerous values. Wrong values
5 could cause unbounded branches in conditionals or unbounded loop executions, or could simply cause wrong
6 calculations and results.

7 There is a special case of pointers or access types. When such a type contains null values, a bound violation and
8 hardware exception can result. When such a type contains plausible but meaningless values, random data reads
9 and writes can collect erroneous data or can destroy data that is in use by another part of the program; when such
10 a type is an access to a subprogram with a plausible (but wrong) value, then either a bad instruction trap may occur
11 or a transfer to an unknown code fragment can occur. All of these scenarios can result in unbounded behaviours.

12 Uninitialized variables are difficult to identify and use for attackers, but can be arbitrarily dangerous in safety
13 situations.

14 **6.29.4 Applicable Language Characteristics**

15 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 16 • Languages that permit variables to be read before they are assigned.

18 **6.29.5 Avoiding the vulnerability or mitigating its effects**

19 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 20 • The general problem of showing that all objects are initialized is intractable; hence developers must
21 carefully structure programs to show that all variables are set before first read on every path throughout the
22 subprogram. Where objects are visible from many modules, it is difficult to determine where the first read
23 occurs, and identify a module that must set the value before that read. When concurrency, interrupts and
24 coroutines are present, it becomes especially imperative to identify where early initialization occurs and to
25 show that the correct order is set via program structure, not by timing, OS precedence, or chance.
- 26 • The simplest method is to initialize each object at elaboration time, or immediately after subprogram
27 execution commences and before any branches. If the subprogram must commence with conditional
28 statements, then the programmer is responsible to show that every variable declared and not initialized
29 earlier is initialized on each branch.
- 30 • Applications can consider defining or reserving fields or portions of the object to only be set when fully
31 initialized.
- 32 • It should be possible to use static analysis tools to show that all objects are set before use in certain
33 specific cases, but as the general problem is intractable, programmers should keep initialization algorithms
34 simple so that they can be analyzed.
- 35 • When declaring and initializing the object together, if the language does not require that the compiler
36 statically verify that the declarative structure and the initialization structure match, use static analysis tools
37 to help detect any mismatches.
- 38 • When setting compound objects, if the language provides mechanisms to set all components together, use
39 those in preference to a sequence of initializations as this helps coverage analysis; otherwise use tools that
40 perform such coverage analysis and document the initialization. Do not perform partial initializations unless
41 there is no choice, and document any deviations from 100% initialization.
- 42 • Where default assignments of multiple components are performed, explicit declaration of the component
43 names and/or ranges helps static analysis and identification of component changes during maintenance.
- 44 • Some languages have named assignments that can be used to build reviewable assignment structures that
45 can be analyzed by the language processor for completeness. Languages with positional notation only can
46 use comments and secondary tools to help show correct assignment.

6.29.6 Implications for standardization

In future standardization activities, the following items should be considered:

- Some languages have ways to determine if modules and regions are elaborated and initialized and to raise exceptions if this does not occur. Languages that do not could consider adding such capabilities.
- Languages could consider setting aside fields in all objects to identify if initialization has occurred, especially for security and safety domains.
- Languages that do not support whole-object initialization could consider adding this capability.

6.30 Wrap-around Error [XYY]

6.30.1 Description of application vulnerability

Wrap-around errors can occur whenever a value is incremented past the maximum or decremented past the minimum value representable in its type and if so specified by the language semantics "wraps around" to either a very small, negative, or undefined value. Using shift operations as a surrogate for multiply or divide may produce a similar error.

6.30.2 Cross reference

CWE:

128. Wrap-around Error
 JSF AV Rules: 164 and 15
 MISRA C 2004: 10.1 to 10.6, 12.8 and 12.11
 MISRA C++ 2008: 2-13-3, 5-0-3 to 5-0-10, and 5-19-1
 CERT C guidelines: INT30-C, INT32-C, and INT34-C

6.30.3 Mechanism of failure

Due to how arithmetic is performed by computers, if a variable is incremented past the maximum value representable in its type, the system may fail to provide an overflow indication to the program. One of the most common processor behaviour is to "wrap" to a very large negative value, or set a condition flag for overflow or underflow, or saturate at the largest representable value.

Shift operations may also produce values that cannot be easily predicted as a result of the different representations of negative integers on various hardware, and, when treating signed quantities, of the differences in behaviour between logical shifts and arithmetic shifts (the particular effect of filling with the sign bit).

Wrap-around often generates an unexpected negative value; this unexpected value may cause a loop to continue for a long time (because the termination condition requires a value greater than some positive value) or an array bounds violation. A wrap-around can sometimes trigger buffer overflows that can be used to execute arbitrary code.

6.30.4 Applicable language characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Languages that do not trigger an exception condition when a wrap-around error occurs.
- Languages that do not fully specify the distinction between arithmetic and logical shifts.

6.30.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 1 • Determine applicable upper and lower bounds for the range of all variables and use language mechanisms
- 2 or static analysis to determine that values are confined to the proper range.
- 3 • Analyze the software using static analysis looking for unexpected consequences of arithmetic operations.
- 4 • Avoid using shift operations as a surrogate for multiplication and division. Most compilers will use the
- 5 correct operation in the appropriate fashion when it is applicable.

6 6.30.6 Implications for standardization

7 In future standardization activities, the following items should be considered:

- 8 • Language standards developers should consider providing facilities to specify either an error, a saturated
- 9 value, or a modulo result when numeric overflow occurs.

10 6.31 Sign Extension Error [XZI]

11 6.31.1 Description of application vulnerability

12 Extending a signed variable that holds a negative value may produce an incorrect result.

13 6.31.2 Cross reference

14 CWE:

15 194. Incorrect Sign Extension

16 MISRA C++ 2008: 5-0-4

17 CERT C guidelines: INT13-C

18 6.31.3 Mechanism of failure

19 Converting a signed data type to a larger data type or pointer can cause unexpected behaviour due to the
 20 extension of the sign bit. A negative data element that is extended with an unsigned extension algorithm will
 21 produce an incorrect result. For instance, this can occur when a signed character is converted to a type short or a
 22 signed integer (32-bit) is converted to an integer type long (64-bit). Sign extension errors can lead to buffer
 23 overflows and other memory based problems. This can occur unexpectedly when moving software designed and
 24 tested on a 32-bit architecture to a 64-bit architecture computer.

25 6.31.4 Applicable language characteristics

26 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 27 • Languages that are weakly typed due to their lack of enforcement of type classifications and interactions.
- 28 • Languages that explicitly or implicitly allow applying unsigned extension operations to signed entities or
- 29 vice-versa.

30 6.31.5 Avoiding the vulnerability or mitigating its effects

31 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 32 • Use a sign extension library, standard function, or appropriate language-specific coding methods to extend
- 33 signed values.
- 34 • Use static analysis tools to help locate situations in which the conversion of variables might have
- 35 unintended consequences.

36 6.31.6 Implications for standardization

37 In future standardization activities, the following items should be considered:

- 1 • Language definitions should define implicit and explicit conversions in a way that prevents alteration of the
2 mathematical value beyond traditional rounding rules.

3 **6.32 Operator Precedence/Order of Evaluation [JCW]**

4 **6.32.1 Description of application vulnerability**

5 Each language provides rules of precedence and associativity, for each expression that operands bind to which
6 operators. These rules are also known as “grouping” or “binding”.

7 Experience and experimental evidence shows that developers can have incorrect beliefs about the relative
8 precedence of many binary operators. See, *Developer beliefs about binary operator precedence*. C Vu, 18(4):14-
9 21, August 2006

10 **6.32.2 Cross reference**

11 JSF AV Rules: 204 and 213
12 MISRA C 2004: 12.1, 12.2, 12.5, 12.6, 13.2, 19.10, 19.12, and 19.13
13 MISRA C++ 2008: 4-5-1, 4-5-2, 4-5-3, 5-0-1, 5-0-2, 5-2-1, 5-3-1, 16-0-6, 16-3-1, and 16-3-2
14 CERT C guidelines: EXP00-C
15 Ada Quaility and Style Guide: 7.1.8 and 7.1.9

16 **6.32.3 Mechanism of failure**

17 In C and C++, the bitwise operators (bitwise logical and bitwise shift) are sometimes thought of by the programmer
18 having similar precedence to arithmetic operations, so just as one might correctly write “ $x - 1 == 0$ ” (“x minus
19 one is equal to zero”), a programmer might erroneously write “ $x \& 1 == 0$ ”, mentally thinking “x anded-with 1 is
20 equal to zero”, whereas the operator precedence rules of C and C++ actually bind the expression as “compute
21 $1==0$, producing ‘false’ interpreted as zero, then bitwise-and the result with x”, producing (a constant) zero,
22 contrary to the programmer’s intent.

23 Examples from an opposite extreme can be found in programs written in APL, which is noteworthy for the absence
24 of *any* distinctions of precedence. One commonly made mistake is to write “ $a * b + c$ ”, intending to produce “a
25 times b plus c”, whereas APL’s uniform right-to-left associativity produces “b plus c, times a”.

26 **6.32.4 Applicable language characteristics**

27 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 28 • Languages whose precedence and associativity rules are sufficiently complex that developers do not
29 remember them.

30 **6.32.5 Avoiding the vulnerability or mitigating its effects**

31 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 32 • Adopt programming guidelines (preferably augmented by static analysis). For example, consider the rules
33 itemized above from JSF C++ [5], CERT C [11] or MISRA C [12].
34 • Use parentheses around binary operator combinations that are known to be a source of error (for example,
35 mixed arithmetic/bitwise and bitwise/relational operator combinations).
36 • Break up complex expressions and use temporary variables to make the order clearer.

37 **6.31.6 Implications for standardization**

38 In future standardization activities, the following items should be considered:

- Language definitions should avoid providing precedence or a particular associativity for operators that are not typically ordered with respect to one another in arithmetic, and instead require full parenthesization to avoid misinterpretation.

6.33 Side-effects and Order of Evaluation [SAM]

6.33.1 Description of application vulnerability

Some programming languages allow subexpressions to cause side-effects (such as assignment, increment, or decrement). For example, some programming languages permit such side-effects, and if, within one expression (such as "`i = v[i++]`"), two or more side-effects modify the same object, undefined behaviour results.

Some languages allow subexpressions to be evaluated in an unspecified ordering, or even removed during optimization. If these subexpressions contain side-effects, then the value of the full expression can be dependent upon the order of evaluation. Furthermore, the objects that are modified by the side-effects can receive values that are dependent upon the order of evaluation.

If a program contains these unspecified or undefined behaviours, testing the program and seeing that it yields the expected results may give the false impression that the expression will always yield the expected result.

6.33.2 Cross reference

JSF AV Rules: 157, 158, 166, 204, 204.1, and 213
 MISRA C 2004: 12.1-12.5
 MISRA C++ 2008: 5-0-1
 CERT C guidelines: EXP10-C, EXP30-C
 Ada Quality and Style Guide: 7.1.8 and 7.1.9

6.33.3 Mechanism of failure

When subexpressions with side effects are used within an expression, the unspecified order of evaluation can result in a program producing different results on different platforms, or even at different times on the same platform. For example, consider

```
a = f(b) + g(b);
```

where `f` and `g` both modify `b`. If `f(b)` is evaluated first, then the `b` used as a parameter to `g(b)` may be a different value than if `g(b)` is performed first. Likewise, if `g(b)` is performed first, `f(b)` may be called with a different value of `b`.

Other examples of unspecified order, or even undefined behaviour, can be manifested, such as

```
a = f(i) + i++;
```

or

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

Parentheses around expressions can assist in removing ambiguity about grouping, but the issues regarding side-effects and order of evaluation are not changed by the presence of parentheses; consider

```
j = i++ * i++;
```

where even if parentheses are placed around the `i++` subexpressions, undefined behaviour still remains. (All examples use the syntax of C or Java for brevity; the effects can be created in any language that allows functions with side-effects in the places where C allows the increment operations.)

1 The unpredictable nature of the calculation means that the program cannot be tested adequately to any degree of
2 confidence. A knowledgeable attacker can take advantage of this characteristic to manipulate data values
3 triggering execution that was not anticipated by the developer.

4 **6.33.4 Applicable language characteristics**

5 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 6 • Languages that permit expressions to contain subexpressions with side effects.
- 7 • Languages whose subexpressions are computed in an unspecified ordering.

9 **6.33.5 Avoiding the vulnerability or mitigating its effects**

10 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 11 • Make use of one or more programming guidelines which (a) prohibit these unspecified or undefined
12 behaviours, and (b) can be enforced by static analysis. (See JSF AV and MISRA rules in Cross reference
13 section [SAM])
- 14 • Keep expressions simple. Complicated code is prone to error and difficult to maintain.

15 **6.33.6 Implications for standardization**

16 In future standardization activities, the following items should be considered:

- 17 • In developing new or revised languages, give consideration to language features that will eliminate or
18 mitigate this vulnerability, such as pure functions.

19 **6.34 Likely Incorrect Expression [KOA]**

20 **6.34.1 Description of application vulnerability**

21 Certain expressions are symptomatic of what is likely to be a mistake made by the programmer. The statement is
22 not wrong, but it is unlikely to be right. The statement may have no effect and effectively is a null statement or may
23 introduce an unintended side-effect. A common example is the use of = in an `if` expression in C where the
24 programmer meant to do an equality test using the `==` operator. Other easily confused operators in C are the
25 logical operators such as `&&` for the bitwise operator `&`, or vice versa. It is legal and possible that the programmer
26 intended to do an assignment within the `if` expression, but due to this being a common error, a programmer doing
27 so would be using a poor programming practice. A less likely occurrence, but still possible is the substitution of `==`
28 for `=` in what is supposed to be an assignment statement, but which effectively becomes a null statement. These
29 mistakes may survive testing only to manifest themselves in deployed code where they may be maliciously
30 exploited.

31 **6.34.2 Cross reference**

32 CWE:

- 33 480. Use of Incorrect Operator
- 34 481. Assigning instead of Comparing
- 35 482. Comparing instead of Assigning
- 36 570. Expression is Always False
- 37 571. Expression is Always True

38 JSF AV Rules: 160 and 166

39 MISRA 2004: 12.3, 12.4, 12.13, 13.1, 13.7, and 14.2

40 MISRA 2008: 0-1-9, 5-0-1, 6-2-1, and 6-5-2

41 CERT C guidelines: MSC02-C and MSC03-C

1 6.34.3 Mechanism of failure

2 Some of the failures are simply a case of programmer carelessness. Substitution of = instead of == in a Boolean
 3 test is easy to do and most C and C++ programmers have made this mistake at one time or another. Other
 4 instances can be the result of intricacies of the language definition that specifies what part of an expression must
 5 be evaluated. For instance, having an assignment expression in a Boolean statement is likely making an
 6 assumption that the complete expression will be executed in all cases. However, this is not always the case as
 7 sometimes the truth-value of the Boolean expression can be determined after only executing some portion of the
 8 expression. For instance:

```
9     if ((a == b) | (c = (d-1)))
```

10 There is no guarantee which of the two subexpressions (a == b) or (c=(d-1)) will be executed first. Should
 11 (a==b) be determined to be true, then there is no need for the subexpression (c=(d-1)) to be executed and as
 12 such, the assignment (c=(d-1)) will not occur.

13 Embedding expressions in other expressions can yield unexpected results. Increment and decrement operators
 14 (++ and --) can also yield unexpected results when mixed into a complex expression.

15 Incorrectly calculated results can lead to a wide variety of erroneous program execution

16 6.34.4 Applicable language characteristics

17 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 18 • All languages are susceptible to likely incorrect expressions.

19 6.34.5 Avoiding the vulnerability or mitigating its effects

20 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 21 • Simplify expressions.
- 22 • Do not use assignment expressions as function parameters. Sometimes the assignment may not be
 23 executed as expected. Instead, perform the assignment before the function call.
- 24 • Do not perform assignments within a Boolean expression. This is likely unintended, but if not, then move
 25 the assignment outside of the Boolean expression for clarity and robustness.
- 26 • On some rare occasions, some statements intentionally do not have side effects and do not cause control
 27 flow to change. These should be annotated through comments and made obvious that they are
 28 intentionally no-ops with a stated reason. If possible, such reliance on null statements should be avoided.
 29 In general, except for those rare instances, all statements should either have a side effect or cause control
 30 flow to change.

31 6.34.6 Implications for standardization

32 In future standardization activities, the following items should be considered:

- 33 • Languages should consider providing warnings for statements that are unlikely to be right such as
 34 statements without side effects. A null (no-op) statement may need to be added to the language for those
 35 rare instances where an intentional null statement is needed. Having a null statement as part of the
 36 language will reduce confusion as to why a statement with no side effects is present in code.
- 37 • Languages should consider not allowing assignments used as function parameters.
- 38 • Languages should consider not allowing assignments within a Boolean expression.
- 39 • Language definitions should avoid situations where easily confused symbols (such as = and ==, or ; and
 40 :, or != and /=) are legal in the same context. For example, = is not generally legal in an if statement in
 41 Java because it does not normally return a boolean value.

1 6.35 Dead and Deactivated Code [XYQ]

2 6.35.1 Description of application vulnerability

3 Dead and Deactivated code (the distinction is addressed in 6.35.3) is code that exists in the executable, but which
4 can never be executed, either because there is no call path that leads to it (for example, a function that is never
5 called), or the path is semantically infeasible (for example, its execution depends on the state of a conditional that
6 can never be achieved).

7 Dead and Deactivated code is undesirable because it indicates the possibility of a coding error and because it may
8 provide a "jump" target for an intrusion.

9 Also covered in this vulnerability is code which is believed to be dead, but which is inadvertently executed.

10 6.35.2 Cross reference

11 CWE:

12 561. Dead Code

13 570. Expression is Always False

14 571. Expression is Always True

15 JSF AV Rules: 127 and 186

16 MISRA C 2004: 14.1 and 2.4

17 MISRA C++ 2008: 0-1-1 to 0-1-10, 2-7-2, and 2-7-3

18 CERT C guidelines: MSC07-C and MSC12-C

19 DO178B/C

20 6.35.3 Mechanism of failure

21 DO-178B defines Dead and Deactivated code as:

- 22 • Dead code – Executable object code (or data) which... cannot be executed (code) or used (data) in an
23 operational configuration of the target computer environment and is not traceable to a system or software
24 requirement.
- 25 • Deactivated code – Executable object code (or data) which by design is either (a) not intended to be
26 executed (code) or used (data), for example, a part of a previously developed software component, or (b) is
27 only executed (code) or used (data) in certain configurations of the target computer environment, for
28 example, code that is enabled by a hardware pin selection or software programmed options.

29 Dead code is code that exists in an application, but which can never be executed, either because there is no call
30 path to the code (for example, a function that is never called) or because the execution path to the code is
31 semantically infeasible, as in

```
32     integer i = 0;
33     if ( i == 0)
34         then fun_a();
35         else fun_b();
```

36 fun_b is dead code, as only fun_a can ever be executed.

37 The presence of dead code is not in itself an error, but begs the question why is it there? Is its presence an
38 indication that the developer believed it to be necessary, but some error means it will never be executed? Or is
39 there a legitimate reason for its presence, for example:

- 40 • Defensive code, only executed as the result of a hardware failure.
- 41 • Code that is part of a library not required in this application.
- 42 • Diagnostic code not executed in the operational environment.

1 Such code may be referred to as “deactivated”. That is, dead code that is there by intent.

2 There is a secondary consideration for dead code in languages that permit overloading of functions and other
3 constructs that use complex name resolution strategies. The developer may believe that some code is not going to
4 be used (deactivated), but its existence in the program means that it appears in the namespace, and may be
5 selected as the best match for some use that was intended to be of an overloading function. That is, although the
6 developer believes it is never going to be used, in practice it is used in preference to the intended function.

7 **6.35.4 Applicable language characteristics**

8 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 9 • Languages that allow code to exist in the executable that can never be executed.

10 **6.35.5 Avoiding the vulnerability or mitigating its effects**

11 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 12 • The developer should endeavour to remove, as a first resort and as far as practical, dead code from an
13 application.
- 14 • When a developer identifies code that is dead because a conditional always evaluates to the same value,
15 this could be indicative of an earlier bug and additional testing may be needed to ascertain why the same
16 value is occurring.
- 17 • The developer should identify any dead code in the application, and provide a justification (if only to
18 themselves) as to why it is there.
- 19 • The developer should also ensure that any code that was expected to be unused is actually recognized as
20 dead code.
- 21 • The developer should apply standard branch coverage measurement tools and ensure by 100% coverage
22 that all branches are neither dead nor deactivated

23 **6.35.6 Implications for standardization**

24 [None]

25 **6.36 Switch Statements and Static Analysis [CLL]**

26 **6.36.1 Description of application vulnerability**

27 Many programming languages provide a construct, such as a `switch` statement, that chooses among multiple
28 alternative control flows based upon the evaluated result of an expression. The use of such constructs may
29 introduce application vulnerabilities if not all possible cases appear within the switch or if control unexpectedly flows
30 from one alternative to another.

31 **6.36.2 Cross reference**

32 JSF AV Rules: 148, 193, 194, 195, and 196
33 MISRA C 2004: 15.2, 15.3, and 15.5
34 MISRA C++ 2008: 6-4-3, 6-4-5, 6-4-6, and 6-4-8
35 CERT C guidelines: MSC01-C
36 Ada Quality and Style Guide: 5.6.1 and 5.6.10

37 **6.36.3 Mechanism of failure**

38 The fundamental challenge when using a `switch` statement is to make sure that all possible cases are, in fact,
39 treated correctly.

6.36.4 Applicable language characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Contain a construct, such as a `switch` statement, that provides a selection among alternative control flows based on the evaluation of an expression.

6.36.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Switch on an expression that has a small number of potential values that can be statically enumerated. In languages that provide them, a variable of an enumerated type is to be preferred because its possible set of values is known statically and is small in number (as compared, for example, to the value set of an integer variable). In languages that don't provide enumerated types, a tightly constrained integer sub-type might be a good alternative. Where it is practical to statically enumerate the switched type, it is preferable to omit the default case, because the static analysis is simplified and because maintainers can better understand the intent of the original programmer. When one must switch on some other form of type, it is necessary to have a default case, preferably to be regarded as a serious error condition.
- Avoid "flowing through" from one case to another. Even if correctly implemented, it is difficult for reviewers and maintainers to distinguish whether the construct was intended or is an error of omission. (Using multiple labels on individual alternatives is not a violation of this recommendation, though.) In cases where flow-through is necessary and intended, an explicitly coded branch may be preferable to clearly mark the intent. Providing comments regarding intention can be helpful to reviewers and maintainers.
- Perform static analysis to determine if all cases are, in fact, covered by the code. (Note that the use of a default case can hamper the effectiveness of static analysis since the tool cannot determine if omitted alternatives were or were not intended for default treatment.)
- Other means of mitigation include manual review, bounds testing, tool analysis, verification techniques, and proofs of correctness.

6.36.6 Implications for standardization

In future standardization activities, the following items should be considered:

- Language specifications could require compilers to ensure that a complete set of alternatives is provided in cases where the value set of the switch variable can be statically determined.

6.37 Demarcation of Control Flow [EOJ]

6.37.1 Description of application vulnerability

Some programming languages explicitly mark the end of an `if` statement or a loop, whereas other languages mark only the end of a block of statements. Languages of the latter category are prone to oversights by the programmer, causing unintended sequences of control flow.

6.37.2 Cross reference

JSF AV Rules: 59 and 192

MISRA C 2004: 14.8, 14.9, 14.10, and 19.5

MISRA C++ 2008: 6-3-1, 6-4-1, 6-4-2, 6-4-3, 6-4-8, 6-5-1, 6-5-6, 6-6-1 to 6-6-5, and 16-0-2

Hatton 18: Control flow – `if` structure

Ada Quality and Style Guide: 3, 5.6.1 through 5.6.10

1 6.37.3 Mechanism of failure

2 Programmers may rely on indentation to determine inclusion of statements within constructs. Testing of the
3 software may not reveal that statements thought to be included in an `if-then`, `if-then-else`, or loops that are
4 not in reality a part of the `if` statement. Moreover, for a nested `if-then-else` statement the programmer may
5 be confused about which `if` statement controls the `else` part directly. This can lead to unexpected results.

6 6.37.4 Applicable language characteristics

7 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 8 • Languages that contain loops and conditional statements that are not explicitly terminated by an “end”
9 construct.

10 6.37.5 Avoiding the vulnerability or mitigating its effects

11 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 12 • Adopt a convention for marking the closing of a construct that can be checked by a tool, to ensure that
13 program structure is apparent.
- 14 • Adopt programming guidelines (preferably augmented by static analysis). For example, consider the rules
15 itemized above from JSF AV, MISRA C, MISRA C++ or Hatton.
- 16 • Other means of assurance might include proofs of correctness, analysis with tools, verification techniques,
17 or other methods.
- 18 • Pretty-printers and syntax-aware editors may be helpful in finding such problems, but sometimes disguise
19 them.
- 20 • Include a final `else` statement at the end of `if-...-else-if` constructs to avoid confusion.
- 21 • Always enclose the body of statements of an `if`, `while`, `for`, or other statements potentially introducing a
22 block of code in braces (“{ }”) or other demarcation indicators appropriate to the language used.

23 6.37.6 Implications for standardization

24 In future standardization activities, the following items should be considered:

- 25 • Specifiers of languages should consider adding a mode that strictly enforces compound conditional and
26 looping constructs with explicit termination, such as “end if” or a closing bracket.
- 27 • Specifiers of languages might consider explicit termination of loops and conditional statements.
- 28 • Specifiers might consider features to terminate named loops and conditionals and determine if the structure
29 as named matches the structure as inferred.

30 6.38 Loop Control Variables [TEX]

31 6.38.1 Description of application vulnerability

32 Many languages support a looping construct whose number of iterations is controlled by the value of a loop control
33 variable. Looping constructs provide a method of specifying an initial value for this loop control variable, a test that
34 terminates the loop and the quantity by which it should be decremented/incremented on each loop iteration.

35 In some languages it is possible to modify the value of the loop control variable within the body of the loop.
36 Experience shows that such value modifications are sometimes overlooked by readers of the source code,
37 resulting in faults being introduced.

38 6.38.2 Cross reference

39 JSF AV Rule: 201

- 1 MISRA C 2004: 13.6
- 2 MISRA C++ 2008: 6-5-1 to 6-5-6

3 **6.38.3 Mechanism of failure**

4 Readers of source code often make assumptions about what has been written. A common assumption is that a
5 loop control variable is a constant since such variables are not usually modified in the body of the associated loop.
6 A reader of the source may incorrectly assume that a loop control variable is not modified in the body of its loop
7 and write (incorrect) code based on this assumption.

8 **6.38.4 Applicable language characteristics**

9 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 10 • Languages that permit a loop control variable to be modified in the body of its associated loop.

11 **6.38.5 Avoiding the vulnerability or mitigating its effects**

12 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 13 • Not modifying a loop control variable in the body of its associated loop body.
- 14 • Some languages, such as C and C++ do not explicitly specify which of the variables appearing in a loop
15 header is the loop control variable. MISRA-C [12] and MISRA C++ [16] have proposed algorithms for
16 deducing which, if any, of these variables is the loop control variable in the programming languages C and
17 C++ (these algorithms could also be applied to other languages that support a C-like for-loop).

18 **6.38.6 Implications for standardization**

19 In future standardization activities, the following items should be considered:

- 20 • Language designers should consider the addition of an identifier type for loop control that cannot be
21 modified by anything other than the loop control construct.

22 **6.39 Off-by-one Error [XZH]**

23 **6.39.1 Description of application vulnerability**

24 A program uses an incorrect maximum or minimum value that is 1 more or 1 less than the correct value. This
25 usually arises from one of a number of situations where the bounds as understood by the developer differ from the
26 design, such as:

- 27 • Confusion between the need for < and <= or > and >= in a test.
- 28 • Confusion as to the index range of an algorithm, such as: beginning an algorithm at 1 when the underlying
29 structure is indexed from 0; beginning an algorithm at 0 when the underlying structure is indexed from 1 (or
30 some other start point); or using the length of a structure as its bound instead of the sentinel values.
- 31 • Failing to allow for storage of a sentinel value, such as the NULL string terminator that is used in the C and
32 C++ programming languages.

33 These issues arise from mistakes in mapping the design into a particular language, in moving between languages
34 (such as between languages where all arrays start at 0 and other languages where arrays start at 1), and when
35 exchanging data between languages with different default array bounds.

36 The issue also can arise in algorithms where relationships exist between components, and the existence of a
37 bounds value changes the conditions of the test.

1 The existence of this possible flaw can also be a serious security hole as it can permit someone to surreptitiously
 2 provide an unused location (such as 0 or the last element) that can be used for undocumented features or hidden
 3 channels.

4 **6.39.2 Cross reference**

5 CWE:
 6 193. Off-by-one Error

7 **6.39.3 Mechanism of failure**

8 An off-by-one error could lead to:

- 9 • an out-of bounds access to an array (buffer overflow),
- 10 • incomplete comparisons or calculation mistakes,
- 11 • a read from the wrong memory location, or
- 12 • an incorrect conditional.

13 Such incorrect accesses can cause cascading errors or references to illegal locations, resulting in potentially
 14 unbounded behaviour.

15 Off-by-one errors are not often exploited in attacks because they are difficult to identify and exploit externally, but
 16 the cascading errors and boundary-condition errors can be severe.

17 **6.39.4 Applicable language characteristics**

18 As this vulnerability arises because of an algorithmic error by the developer, it can in principle arise in any
 19 language; however, it is most likely to occur when:

- 20 • The language relies on the developer having implicit knowledge of structure start and end indices (for
 21 example, knowing whether arrays start at 0 or 1 – or indeed some other value).
- 22 • Where the language relies upon explicit bounds values to terminate variable length arrays.

23 **6.39.5 Avoiding the vulnerability or mitigating its effects**

24 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 25 • A systematic development process, use of development/analysis tools and thorough testing are all
 26 common ways of preventing errors, and in this case, off-by-one errors.
- 27 • Where references are being made to structure indices and the languages provide ways to specify the
 28 whole structure or the starting and ending indices explicitly (for example., Ada provides xxx'First and
 29 xxx'Last for each dimension), these should be used always. Where the language doesn't provide these,
 30 constants can be declared and used in preference to numeric literals.
- 31 • Where the language doesn't encapsulate variable length arrays, encapsulation should be provided through
 32 library objects and a coding standard developed that requires such arrays to only be used via those library
 33 objects, so the developer does not need to be explicitly concerned with managing bounds values.
 34

35 **6.39.6 Implications for standardization**

36 In future standardization activities, the following items should be considered:

- 37 • Languages should provide encapsulations for arrays that:
 - 38 ○ Prevent the need for the developer to be concerned with explicit bounds values.
 - 39 ○ Provide the developer with symbolic access to the array start, end and iterators.

1 6.40 Structured Programming [EWD]

2 6.40.1 Description of application vulnerability

3 Programs that have a convoluted control structure are likely to be more difficult to be human readable, less
4 understandable, harder to maintain, more difficult to modify, harder to statically analyze, more difficult to match the
5 allocation and release of resources, and more likely to be incorrect.

7 6.40.2 Cross reference

8 JSF AV Rules: 20, 113, 189, 190, and 191
9 MISRA C 2004: 14.4, 14.5, and 20.7
10 MISRA C++ 2008: 6-6-1, 6-6-2, 6-6-3, and 17-0-5
11 CERT C guidelines: SIG32-C
12 Ada Quality and Style Guide: 3, 4, 5.4, 5.6, and 5.7
13

14 6.40.3 Mechanism of failure

15 Lack of structured programming can lead to:

- 16 • Memory or resource leaks.
- 17 • Error prone maintenance.
- 18 • Design that is difficult or impossible to validate.
- 19 • Source code that is difficult or impossible to statically analyze.

20 6.40.4 Applicable language characteristics

21 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 22 • Languages that allow leaving a loop without consideration for the loop control.
- 23 • Languages that allow local jumps (`goto` statement).
- 24 • Languages that allow non-local jumps (`set jmp/long jmp` in the C programming language).
- 25 • Languages that support multiple entry and exit points from a function, procedure, subroutine or method.

26 6.40.5 Avoiding the vulnerability or mitigating its effects

27 Use only those features of the programming language that enforce a logical structure on the program. The
28 program flow follows a simple hierarchical model that employs looping constructs such as `for`, `repeat`, `do`, and
29 `while`.

30 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 31 • Avoid using language features such as `goto`.
- 32 • Avoid using language features such as `continue` and `break` in the middle of loops.
- 33 • Avoid using language features that transfer control of the program flow via a jump.
- 34 • Avoid multiple exit points to a function/procedure/method/subroutine.
- 35 • Avoid multiple entry points to a function/procedure/method/subroutine.

36 6.40.6 Implications for standardization

37 In future standardization activities, the following items should be considered:

- 38 • Languages should support and favor structured programming through their constructs to the extent
39 possible.

1 6.41 Passing Parameters and Return Values [CSJ]

2 6.41.1 Description of application vulnerability

3 Nearly every procedural language provides some method of process abstraction permitting decomposition of the
 4 flow of control into routines, functions, subprograms, or methods. (For the purpose of this description, the term
 5 subprogram will be used.) To have any effect on the computation, the subprogram must change data visible to the
 6 calling program. It can do this by changing the value of a non-local variable, changing the value of a parameter, or,
 7 in the case of a function, providing a return value. Because different languages use different mechanisms with
 8 different semantics for passing parameters, a programmer using an unfamiliar language may obtain unexpected
 9 results.

10 6.41.2 Cross reference

11 JSF AV Rules: 116, 117, and 118
 12 MISRA C 2004: 16.1, 16.2, 16.3, 16.4, 16.5, 16.6, 16.7, and 16.9
 13 MISRA C++ 2008: 0-3-2, 7-1-2, 8-4-1, 8-4-2, 8-4-3, and 8-4-4
 14 CERT C guidelines: EXP12-C and DCL33-C
 15 Ada Quality and Style Guide: 5.2 and 8.3

16 6.41.3 Mechanism of failure

17 The mechanisms for parameter passing include: *call by reference*, *call by copy*, and *call by name*. The last is so
 18 specialized and supported by so few programming languages that it will not be treated in this description.

19 In call by reference, the calling program passes the addresses of the arguments to the called subprogram. When
 20 the subprogram references the corresponding formal parameter, it is actually sharing data with the calling program.
 21 If the subprogram changes a formal parameter, then the corresponding actual argument is also changed. If the
 22 actual argument is an expression or a constant, then the address of a temporary location is passed to the
 23 subprogram; this may be an error in some languages.

24 In call by copy, the called subprogram does not share data with the calling program. Instead, formal parameters act
 25 as local variables. Values are passed between the actual arguments and the formal parameters by copying. Some
 26 languages may control changes to formal parameters based on labels such as *in*, *out*, or *inout*. There are three
 27 cases to consider: *call by value* for *in* parameters; *call by result* for *out* parameters and function return values;
 28 and *call by value-result* for *inout* parameters. For call by value, the calling program evaluates the actual
 29 arguments and copies the result to the corresponding formal parameters that are then treated as local variables by
 30 the subprogram. For call by result, the values of the locals corresponding to formal parameters are copied to the
 31 corresponding actual arguments. For call by value-result, the values are copied in from the actual arguments at the
 32 beginning of the subprogram's execution and back out to the actual arguments at its termination.

33 The obvious disadvantage of call by copy is that extra copy operations are needed and execution time is required
 34 to produce the copies. Particularly if parameters represent sizable objects, such as large arrays, the cost of call by
 35 copy can be high. For this reason, many languages also provide the call by reference mechanism. The
 36 disadvantage of call by reference is that the calling program cannot be assured that the subprogram hasn't
 37 changed data that was intended to be unchanged. For example, if an array is passed by reference to a subprogram
 38 intended to sum its elements, the subprogram could also change the values of one or more elements of the array.
 39 However, some languages enforce the subprogram's access to the shared data based on the labeling of actual
 40 arguments with modes—such as *in*, *out*, or *inout* or by constant pointers.

41 Another problem with call by reference is unintended aliasing. It is possible that the address of one actual argument
 42 is the same as another actual argument or that two arguments overlap in storage. A subprogram, assuming the two
 43 formal parameters to be distinct, may treat them inappropriately. For example, if one codes a subprogram to swap
 44 two values using the exclusive-or method, then a call to `swap(x, x)` will zero the value of `x`. Aliasing can also
 45 occur between arguments and non-local objects. For example, if a subprogram modifies a non-local object as a
 46 side-effect of its execution, referencing that object by a formal parameter will result in aliasing and, possibly,
 47 unintended results.

1 Some languages provide only simple mechanisms for passing data to subprograms, leaving it to the programmer to
 2 synthesize appropriate mechanisms. Often, the only available mechanism is to use call by copy to pass small
 3 scalar values or pointer values containing addresses of data structures. Of course, the latter amounts to using call
 4 by reference with no checking by the language processor. In such cases, subprograms can pass back pointers to
 5 anything whatsoever, including data that is corrupted or absent.

6 Some languages use call by copy for small objects, such as scalars, and call by reference for large objects, such
 7 as arrays. The choice of mechanism may even be implementation-defined. Because the two mechanisms produce
 8 different results in the presence of aliasing, it is very important to avoid aliasing.

9 An additional problem may occur if the called subprogram fails to assign a value to a formal parameter that the
 10 caller expects as an output from the subprogram. In the case of call by reference, the result may be an uninitialized
 11 variable in the calling program. In the case of call by copy, the result may be that a legitimate initialization value
 12 provided by the caller is overwritten by an uninitialized value because the called program did not make an
 13 assignment to the parameter. This error may be difficult to detect through review because the failure to initialize is
 14 hidden in the subprogram.

15 An additional complication with subprograms occurs when one or more of the arguments are expressions. In such
 16 cases, the evaluation of one argument might have side-effects that result in a change to the value of another or
 17 unintended aliasing. Implementation choices regarding order of evaluation could affect the result of the
 18 computation. This particular problem is described in Side-effects and Order of Evaluation section [SAM].

19 **6.41.4 Applicable language characteristics**

20 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 21 • Languages that provide mechanisms for defining subprograms where the data passes between the calling
 22 program and the subprogram via parameters and return values. This includes methods in many popular
 23 object-oriented languages.

24 **6.41.5 Avoiding the vulnerability or mitigating its effects**

25 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 26 • Use available mechanisms to label parameters as constants or with modes like `in`, `out`, or `inout`.
- 27 • When a choice of mechanisms is available, pass small simple objects using call by copy.
- 28 • When a choice of mechanisms is available and the computational cost of copying is tolerable, pass larger
 29 objects using call by copy.
- 30 • When the choice of language or the computational cost of copying forbids using call by copy, then take
 31 safeguards to prevent aliasing:
 - 32 ○ Minimize side-effects of subprograms on non-local objects; when side-effects are coded, ensure
 33 that the affected non-local objects are not passed as parameters using call by reference.
 - 34 ○ To avoid unintentional aliasing, avoid using expressions or functions as actual arguments; instead
 35 assign the result of the expression to a temporary local and pass the local.
 - 36 ○ Utilize tooling or other forms of analysis to ensure that non-obvious instances of aliasing are
 37 absent.
 - 38 ○ Perform reviews or analysis to determine that called subprograms fulfil their responsibilities to
 39 assign values to all output parameters.

40 **6.41.6 Implications for standardization**

41 In future standardization activities, the following items should be considered:

- 42 • Programming language specifications could provide labels—such as `in`, `out`, and `inout`—that control the
 43 subprogram's access to its formal parameters, and enforce the access.

1 6.42 Dangling References to Stack Frames [DCM]

2 6.42.1 Description of application vulnerability

3 Many languages allow treating the address of a local variable as a value stored in other variables. Examples are
 4 the application of the address operator in C or C++, or of the 'Access or 'Address attributes in Ada. In some
 5 languages, this facility is also used to model the call-by-reference mechanism by passing the address of the actual
 6 parameter by-value. An obvious safety requirement is that the stored address shall not be used after the lifetime of
 7 the local variable has expired. Technically, the stack frame, in which the local variable lived, has been popped and
 8 memory may have been reused for a subsequent call. Therefore, the invalidity of the stored address is very difficult
 9 to decide. This situation can be described as a "dangling reference to the stack".

10 6.42.2 Cross reference

11 CWE:

12 562. Return of Stack Variable Address

13 JSF AV Rule: 173

14 MISRA C 2004: 17.6 and 21.1

15 MISRA C++ 2008: 0-3-1, 7-5-1, 7-5-2, and 7-5-3

16 CERT C guidelines: EXP35-C and DCL30-C

17 Ada Quality and Style Guide: 7.6.7, 7.6.8, and 10.7.6

18 6.42.3 Mechanism of failure

19 The consequences of dangling references to the stack come in two variants: a deterministically predictable variant,
 20 which therefore can be exploited, and an intermittent, non-deterministic variant, which is next to impossible to elicit
 21 during testing. The following code sample illustrates the two variants; the behaviour is not language-specific:

```

22     struct s { ... };
23     typedef struct s array_type[1000];
24     array_type* ptr;
25     array_type* F()
26     {
27         struct s Arr[1000];
28         ptr = &Arr;      // Risk of variant 1;
29         return &Arr;    // Risk of variant 2;
30     }
31     ...
32     struct s secret;
33     array_type* ptr2;
34     ptr2 = F();
35     secret = (*ptr2)[10]; // Fault of variant 2
36     ...
37     secret = (*ptr)[10]; // Fault of variant 1
  
```

38 The risk of variant 1 is the assignment of the address of `Arr` to a pointer variable that survives the lifetime of `Arr`.
 39 The fault is the subsequent use of the dangling reference to the stack, which references memory since altered by
 40 other calls and possibly validly owned by other routines. As part of a call-back, the fault allows systematic
 41 examination of portions of the stack contents without triggering an array-bounds-checking violation. Thus, this
 42 vulnerability is easily exploitable. As a fault, the effects can be most astounding, as memory gets corrupted by
 43 completely unrelated code portions. (A life-time check as part of pointer assignment can prevent the risk. In many
 44 cases, such as the situations above, the check is statically decidable by a compiler. However, for the general case,
 45 a dynamic check is needed to ensure that the copied pointer value lives no longer than the designated object.)

46 The risk of variant 2 is an idiom "seen in the wild" to return the address of a local variable to avoid an expensive
 47 copy of a function result, as long as it is consumed before the next routine call occurs. The idiom is based on the ill-
 48 founded assumption that the stack will not be affected by anything until this next call is issued. The assumption is

1 false, however, if an interrupt occurs and interrupt handling employs a strategy called “stack stealing”, that is, using
 2 the current stack to satisfy its memory requirements. Thus, the value of `ARR` can be overwritten before it can be
 3 retrieved after the call on `F`. As this fault will only occur if the interrupt arrives after the call has returned but before
 4 the returned result is consumed, the fault is highly intermittent and next to impossible to re-create during testing.
 5 Thus, it is unlikely to be exploitable, but also exceedingly hard to find by testing. It can begin to occur after a
 6 completely unrelated interrupt handler has been coded or altered. Only static analysis can relatively easily detect
 7 the danger (unless the code combines it with risks of variant 1). Some compilers issue warnings for this situation;
 8 such warnings need to be heeded, and some forms of static analysis are effective in identifying such problems.

9 **6.42.4 Applicable language characteristics**

10 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 11 • The address of a local entity (or formal parameter) of a routine can be obtained and stored in a variable or
 12 can be returned by this routine as a result.
- 13 • No check is made that the lifetime of the variable receiving the address is no larger than the lifetime of the
 14 designated entity.

15 **6.42.5 Avoiding the vulnerability or mitigating its effects**

16 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 17 • Do not use the address of locally declared entities as storable, assignable or returnable value (except
 18 where idioms of the language make it unavoidable).
- 19 • Where unavoidable, ensure that the lifetime of the variable containing the address is completely enclosed
 20 by the lifetime of the designated object.
- 21 • Never return the address of a local variable as the result of a function call.

22 **6.42.6 Implications for standardization**

23 In future standardization activities, the following items should be considered:

- 24 • Do not provide means to obtain the address of a locally declared entity as a storable value; or
- 25 • Define implicit checks to implement the assurance of enclosed lifetime expressed in subsection 5 of this
 26 vulnerability. Note that, in many cases, the check is statically decidable, for example, when the address of
 27 a local entity is taken as part of a return statement or expression.

28 **6.43 Subprogram Signature Mismatch [OTR]**

29 **6.43.1 Description of application vulnerability**

30 If a subprogram is called with a different number of parameters than it expects, or with parameters of different types
 31 than it expects, then the results will be incorrect. Depending on the language, the operating environment, and the
 32 implementation, the error might be as benign as a diagnostic message or as extreme as a program continuing to
 33 execute with a corrupted stack. The possibility of a corrupted stack provides opportunities for penetration.

34 **6.43.2 Cross reference**

35 CWE:

- 36 628. Function Call with Incorrectly Specified Arguments
- 37 686. Function Call with Incorrect Argument Type
- 38 683. Function Call with Incorrect Order of Arguments

39 JSF AV Rule: 108

40 MISRA C 2004: 8.1, 8.2, 8.3, 16.1, 16.3, 16.4, 16.5, 16.6

41 MISRA C++ 2008: 0-3-2, 3-2-1, 3-2-2, 3-2-3, 3-2-4, 3-3-1, 3-9-1, 8-3-1, 8-4-1, and 8-4-2

42 CERT C guidelines: DCL31-C, and DCL35-C

1

2 6.43.3 Mechanism of failure

3 When a subprogram is called, the actual arguments of the call are pushed on to the execution stack. When the
4 subprogram terminates, the formal parameters are popped off the stack. If the number and type of the actual
5 arguments do not match the number and type of the formal parameters, then the push and the pop will not be
6 commensurable and the stack will be corrupted. Stack corruption can lead to unpredictable execution of the
7 program and can provide opportunities for execution of unintended or malicious code.

8 The compilation systems for many languages and implementations can check to ensure that the list of actual
9 parameters and any expected return match the declared set of formal parameters and return value (the
10 *subprogram signature*) in both number and type. (In some cases, programmers should observe a set of
11 conventions to ensure that this is true.) However, when the call is being made to an externally compiled
12 subprogram, an object-code library, or a module compiled in a different language, the programmer must take
13 additional steps to ensure a match between the expectations of the caller and the called subprogram.

14 6.43.4 Applicable language characteristics

15 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 16 • Languages that do not require their implementations to ensure that the number and types of actual
17 arguments are equal to the number and types of the formal parameters.
- 18 • Implementations that permit programs to call subprograms that have been externally compiled (without a
19 means to check for a matching subprogram signature), subprograms in object code libraries, and any
20 subprograms compiled in other languages.

21 6.43.5 Avoiding the vulnerability or mitigating its effects

22 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 23 • Take advantage of any mechanism provided by the language to ensure that subprogram signatures
24 match.
- 25 • Avoid any language features that permit variable numbers of actual arguments without a method of
26 enforcing a match for any instance of a subprogram call.
- 27 • Take advantage of any language or implementation feature that would guarantee matching the
28 subprogram signature in linking to other languages or to separately compiled modules.
- 29 • Intensively review subprogram calls where the match is not guaranteed by tooling.

30 6.43.6 Implications for standardization

31 In future standardization activities, the following items should be considered:

- 32 • Language specifiers could ensure that the signatures of subprograms match within a single compilation
33 unit and could provide features for asserting and checking the match with externally compiled
34 subprograms.

35 6.44 Recursion [GDL]

36 6.44.1 Description of application vulnerability

37 Recursion is an elegant mathematical mechanism for defining the values of some functions. It is tempting to write
38 code that mirrors the mathematics. However, the use of recursion in a computer can have a profound effect on the
39 consumption of finite resources, leading to denial of service.

1 **6.44.2 Cross reference**

2 CWE:

- 3 674. Uncontrolled Recursion
- 4 JSF AV Rule: 119
- 5 MISRA C 2004: 16.2
- 6 MISRA C++ 2008: 7-5-4
- 7 CERT C guidelines: MEM05-C
- 8 Ada Quality and Style Guide: 5.6.6

9 **6.44.3 Mechanism of failure**

10 Recursion provides for the economical definition of some mathematical functions. However, economical definition
 11 and economical calculation are two different subjects. It is tempting to calculate the value of a recursive function
 12 using recursive subprograms because the expression in the programming language is straightforward and easy to
 13 understand. However, the impact on finite computing resources can be profound. Each invocation of a recursive
 14 subprogram may result in the creation of a new stack frame, complete with local variables. If stack space is limited
 15 and the calculation of some values will lead to an exhaustion of resources resulting in the program terminating.

16 In calculating the values of mathematical functions the use of recursion in a program is usually obvious, but this is
 17 not true in the general case. For example, finalization of a computing context after treating an error condition might
 18 result in recursion (such as attempting to "clean up" by closing a file after an error was encountered in closing the
 19 same file). Although such situations may have other problems, they typically do not result in exhaustion of
 20 resources but may otherwise result in a denial of service.

21 **6.44.4 Applicable language characteristics**

22 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 23 • Any language that permits the recursive invocation of subprograms.

24 **6.44.5 Avoiding the vulnerability or mitigating its effects**

25 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 26 • Minimize the use of recursion.
- 27 • Converting recursive calculations to the corresponding iterative calculation. In principle, any recursive
 28 calculation can be remodeled as an iterative calculation which will have a smaller impact on some
 29 computing resources but which may be harder for a human to comprehend. The cost to human
 30 understanding must be weighed against the practical limits of computing resource.
- 31 • In cases where the depth of recursion can be shown to be statically bounded by a tolerable number, then
 32 recursion may be acceptable, but should be documented for the use of maintainers.

33 It should be noted that some languages or implementations provide special (more economical) treatment of a form
 34 of recursion known as *tail-recursion*. In this case, the impact on computing economy is reduced. When using such
 35 a language, tail recursion may be preferred to an iterative calculation.

36 **6.44.6 Implications for standardization**

37 [None]

1 **6.45 Returning Error Status [NZN]**

2 **6.45.1 Description of application vulnerability**

3 Unpredicted error conditions—perhaps from hardware (such as an I/O device error), perhaps from software (such
4 as heap exhaustion)—sometimes arise during the execution of code. Programming languages provide a
5 surprisingly wide variety of mechanisms to deal with such errors. The choice of a mechanism that doesn't match
6 the programming language can lead to errors in the execution of the software or unexpected termination of the
7 program. This could lead to a simple decrease in the robustness of a program or it could be exploited in a denial of
8 service attack.

9 **6.45.2 Cross reference**

10 JSF AV Rules: 115 and 208
11 MISRA C 2004: 16.10
12 MISRA C++ 2008: 15-3-2 and 19-3-1
13 CERT C guidelines: DCL09-C, ERR00-C, and ERR02-C

14 **6.45.3 Mechanism of failure**

15 Even in the best-written programs, error conditions sometimes arise. Some errors occur because of defects in the
16 software itself, but some result from external conditions in hardware, such as errors in I/O devices, or in the
17 software system, such as exhaustion of heap space. If left untreated, the effect of the error might result in
18 termination of the program or continuation of the program with incorrect results. To deal with the situation,
19 designers of programming languages have equipped their languages with different mechanisms to detect and treat
20 such errors. These mechanisms are typically intended to be used in specific programming idioms. However, the
21 mechanisms differ among languages. A programmer expert in one language might mistakenly use an inappropriate
22 idiom when programming in a different language with the result that some errors are left untreated, leading to
23 termination or incorrect results. Attackers can exploit such weaknesses in denial of service attacks.

24 In general, languages make no distinction between dealing with programming errors (like an access to protected
25 memory), unexpected hardware errors (like device error), expected but unusual conditions (like end of file), and
26 even usual conditions that fail to provide the typical result (like an unsuccessful search). This description will use
27 the term "error" to apply to all of the above. The description applies equally to error conditions that are detected via
28 hardware mechanisms and error conditions that are detected via software during execution of a subprogram (such
29 as an inappropriate parameter value).

30 **6.45.4 Applicable language characteristics**

31 Different programming languages provide remarkably different mechanisms for treating errors. In languages that
32 provide a number of error detection and treatment mechanisms, it becomes a design issue to match the
33 mechanism to the condition. This section will describe the mechanisms that are provided in widely used languages.

34 The simplest case is the set of languages that provide no special mechanism for the notification and treatment of
35 unusual conditions. In such languages, error conditions are signaled by the value of an auxiliary status variable,
36 sometimes a subprogram parameter. The programming language C standard library functions use a variant of this
37 approach; the error status is provided as the return value and sometimes in an additional global error value.
38 Obviously, in such languages, it is imperative to check and act upon the status variable after every call to a
39 subprogram that might provide an error indication. If error conditions can occur in an asynchronous manner, it is
40 necessary to provide means to check for errors in a systematic and periodic manner.

41 Some languages permit the passing of a label parameter. If an error is encountered, the subprogram returns to the
42 indicated label rather than to the point at which it was called. Similarly some languages accept the name of a
43 subprogram to be used to handle errors. In either case, it is imperative to provide labeled code or a subprogram to
44 deal with all possible error situations.

The approaches described above have the disadvantage that error checking must be provided at every call to a subprogram. This can clutter the code immensely to deal with situations that may occur rarely. For this reason, some languages provide an exception mechanism that automatically transfers control when an error is encountered. This has the potential advantage of allowing error treatment to be factored into distinct error handlers, leaving the main execution path to deal with the usual results. The disadvantages, of course, are that the language design is complicated and the programmer must deal with the conceptually more complex problem of providing error handlers that are removed from the immediate context of a specific call to a subprogram. Furthermore, different languages provide exception-handling mechanisms that differ in the manner in which various design issues are treated:

- How is the occurrence of an exception bound to a particular handler?
- What happens when no handler is local to an exception occurrence? Is the exception propagated in some manner or is it lost?
- What happens after an exception handler executes? Is control returned to the point before the call or after the call, or is the calling routine terminated in some way? If the calling routine is terminated, is there some provision for finalization, such as closing files or releasing resources?
- Are programmers permitted to define additional exceptions?
- Does the language provide default handlers for some exceptions or must the programmer explicitly provide for all of them?
- Can predefined exceptions be raised explicitly?
- Under what circumstances can error checking be disabled?

6.45.5 Avoiding the vulnerability or mitigating its effects

Given the variety of error handling mechanisms, it is difficult to write general guidelines. However, dealing with exception handlers can stress the capability of many static analysis tools and can, in some cases, reduce the effectiveness of their analysis. Therefore, for situations where the highest of reliability is required, the application should be designed so that exception handling is not used at all. In the more general case, exception-handling mechanisms should be reserved for truly unexpected situations and other situations (possibly hardware arithmetic overflow) where no other mechanism is available. Situations which are merely unusual, like end of file, should be treated by explicit testing—either prior to the call which might raise the error or immediately afterward.

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Checking error return values or auxiliary status variables following a call to a subprogram is mandatory unless it can be demonstrated that the error condition is impossible.
- In dealing with languages where untreated exceptions can be lost (for example, an exception that goes untreated within an Ada task), it is mandatory to deal with the exception in the local context before it is lost.
- When execution within a particular context is abandoned due to an exception, it is important to finalize the context by closing open files, releasing resources and restoring any invariants associated with the context.
- It is often not appropriate to repair an error condition and retry the operation. In such cases, one often treats a symptom but not the underlying problem. It is usually a better solution to finalize and terminate the current context and retreat to a context where the situation is known.
- Error checking provided by the language, the software system, or the hardware should never be disabled in the absence of a conclusive analysis that the error condition is rendered impossible.
- Because of the complexity of error handling, careful review of all error handling mechanisms is appropriate.
- In applications with the highest requirements for reliability, defense-in-depth approaches are often appropriate, for example, checking and handling errors thought to be impossible.

6.45.6 Implications for standardization

In future standardization activities, the following items should be considered:

- A standardized set of mechanisms for detecting and treating error conditions should be developed so that all languages to the extent possible could use them. This does not mean that all languages should use the same mechanisms as there should be a variety (for example, label parameters, auxiliary status variables), but each of the mechanisms should be standardized.

1 **6.46 Termination Strategy [REU]**

2 **6.46.1 Description of application vulnerability**

3 Expectations that a system will be dependable are based on the confidence that the system will operate as
 4 expected and not fail in normal use. The dependability of a system and its fault tolerance can be measured
 5 through the component part's reliability, availability, safety and security. Reliability is the ability of a system or
 6 component to perform its required functions under stated conditions for a specified period of time [IEEE 1990
 7 glossary]. Availability is how timely and reliable the system is to its intended users. Both of these factors matter
 8 highly in systems used for safety and security. In spite of the best intentions, systems may encounter a failure,
 9 either from internally poorly written software or external forces such as power outages/variations, floods, or other
 10 natural disasters. The reaction to a fault can affect the performance of a system and in particular, the safety and
 11 security of the system and its users.

12 When the software does not terminate in the planned mechanism, safety or security is compromised, as failing in
 13 an unspecified way interferes with the alternative recovery features. In safety-related systems the results can be
 14 catastrophic: for other systems the result can mean failure of the complete system

15 **6.46.2 Cross reference**

16 JSF AV Rule: 24
 17 MISRA C 2004: 20.11
 18 MISRA C++ 2008: 0-3-2, 15-5-2, 15-5-3, and 18-0-3
 19 CERT C guidelines: ERR04-C, ERR06-C and ENV32-C
 20 Ada Quality and Style Guide: 5.8 and 7.5

21 **6.46.3 Mechanism of failure**

22 The reaction to a fault in a system can depend on the criticality of the part in which the fault originates. When a
 23 program consists of several tasks, each task may be critical, or not. If a task is critical, it may or may not be
 24 restartable by the rest of the program. Ideally, a task that detects a fault within itself should be able to halt leaving
 25 its resources available for use by the rest of the program, halt clearing away its resources, or halt the entire
 26 program. The latency of task termination and whether tasks can ignore termination signals should be clearly
 27 specified. Having inconsistent reactions to a fault can potentially be a vulnerability.

28 When a fault is detected, there are many ways in which a system can react. The quickest and most noticeable way
 29 is to fail hard, also known as fail fast or fail stop. The reaction to a detected fault is to immediately halt the system.
 30 Alternatively, the reaction to a detected fault could be to fail soft. The system would keep working with the faults
 31 present, but the performance of the system would be degraded. Systems used in a high availability environment
 32 such as telephone switching centers, e-commerce, or other "always available" applications would likely use a fail
 33 soft approach. What is actually done in a fail soft approach can vary depending on whether the system is used for
 34 safety critical or security critical purposes. For fail-safe systems, such as flight controllers, traffic signals, or
 35 medical monitoring systems, there would be no effort to meet normal operational requirements, but rather to limit
 36 the damage or danger caused by the fault. A system that fails securely, such as cryptologic systems, would
 37 maintain maximum security when a fault is detected, possibly through a denial of service.

38 **6.46.4 Applicable language characteristics**

39 This vulnerability description is intended to be applicable to all languages.

40 **6.46.5 Avoiding the vulnerability or mitigating its effects**

41 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 42 • A strategy for fault handling should be decided. Consistency in fault handling should be the same with
 43 respect to critically similar parts.
- 44 • A multi-tiered approach of fault prevention, fault detection and fault reaction should be used.

- 1 • System-defined components that assist in uniformity of fault handling should be used when available. For
2 one example, designing a "runtime constraint handler" (as described in ISO/IEC TR 24731-1 [13]) permits
3 the application to intercept various erroneous situations and perform one consistent response, such as
4 flushing a previous transaction and re-starting at the next one.
- 5 • When there are multiple tasks, a fault-handling policy should be specified whereby a task may
6 o halt, and keep its resources available for other tasks (perhaps permitting restarting of the faulting
7 task)
- 8 o halt, and remove its resources (perhaps to allow other tasks to use the resources so freed, or to
9 allow a recreation of the task)
- 10 o halt, and signal the rest of the program to likewise halt.

11 **6.46.6 Implications for standardization**

12 In future standardization activities, the following items should be considered:

- 13 • Languages should consider providing a means to perform fault handling. Terminology and the means
14 should be coordinated with other languages.

15 **6.47 Extra Intrinsic [LRM]**

16 **6.47.1 Description of application vulnerability**

17 Most languages define intrinsic procedures, which are easily available, or always "simply available", to any
18 translation unit. If a translator extends the set of intrinsics beyond those defined by the standard, and the standard
19 specifies that intrinsics are selected before procedures of the same signature defined by the application, a different
20 procedure may be unexpectedly used when switching between translators.

21 **6.47.2 Cross reference**

22 [None]

23 **6.47.3 Mechanism of failure**

24 Most standard programming languages define a set of intrinsic procedures which may be used in any application.
25 Some language standards allow a translator to extend this set of intrinsic procedures. Some language standards
26 specify that intrinsic procedures are selected ahead of an application procedure of the same signature. This may
27 cause a different procedure to be used when switching between translators.

28 For example, most languages provide a routine to calculate the square root of a number, usually named `sqrt()`. If
29 a translator also provided, as an extension, a cube root routine, say named `cbrt()`, that extension may override
30 an application defined procedure of the same signature. If the two different `cbrt()` routines chose different branch
31 cuts when applied to complex arguments, the application could unpredictably go wrong.

32 If the language standard specifies that application defined procedures are selected ahead of intrinsic procedures of
33 the same signature, the use of the wrong procedure may mask a linking error.

34 **6.47.4 Applicable language characteristics**

35 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 36 • Any language where translators may extend the set of intrinsic procedures and where intrinsic procedures
37 are selected ahead of application defined (or external library defined) procedures of the same signature.

38 **6.47.5 Avoiding the vulnerability or mitigating its effects**

39 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 1 • Use whatever language features are available to mark a procedure as language defined or application
- 2 defined.
- 3 • Be aware of the documentation for every translator in use and avoid using procedure signatures matching
- 4 those defined by the translator as extending the standard set.

5 **6.47.6 Implications for standardization**

6 In future standardization activities, the following items should be considered:

- 7 • Clearly state whether translators can extend the set of intrinsic procedures or not.
- 8 • Clearly state what the precedence is for resolving collisions.
- 9 • Clearly provide ways to mark a procedure signature as being the intrinsic or an application provided
- 10 procedure.
- 11 • Require that a diagnostic is issued when an application procedure matches the signature of an intrinsic
- 12 procedure.

13 **6.48 Type-breaking Reinterpretation of Data [AMV]**

14 **6.48.1 Description of application vulnerability**

15 In most cases, objects in programs are assigned locations in processor storage to hold their value. If the same
 16 storage space is assigned to more than one object—either statically or temporarily—then a change in the value of
 17 one object will have an effect on the value of the other. Furthermore, if the representation of the value of an object
 18 is reinterpreted as being the representation of the value of an object with a different type, unexpected results may
 19 occur.

20 **6.48.2 Cross reference**

21 JSF AV Rules 153 and 183
 22 MISRA 2004: 18.2, 18.3, and 18.4
 23 MISRA C++ 2008: 4-5-1 to 4-5-3, 4-10-1, 4-10-2, and 5-0-3 to 5-0-9
 24 CERT C guidelines: MEM08-C
 25 Ada Quality and Style Guide: 7.6.7 and 7.6.8

26 **6.48.3 Mechanism of failure**

27 Sometimes there is a legitimate need for applications to place different interpretations upon the same stored
 28 representation of data. The most fundamental example is a program loader that treats a binary image of a program
 29 as data by loading it, and then treats it as a program by invoking it. Most programming languages permit type-
 30 breaking reinterpretation of data, however, some offer less error prone alternatives for commonly encountered
 31 situations.

32 Type-breaking reinterpretation of representation presents obstacles to human understanding of the code, the ability
 33 of tools to perform effective static analysis, and the ability of code optimizers to do their job.

34 Examples include:

- 35 • Providing alternative mappings of objects into blocks of storage performed either statically (such as Fortran
- 36 `COMMON`) or dynamically (such as pointers).
- 37 • Union types, particularly unions that do not have a discriminant stored as part of the data structure.
- 38 • Operations that permit a stored value to be interpreted as a different type (such as treating the
- 39 representation of a pointer as an integer).

40 In all of these cases accessing the value of an object may produce an unanticipated result.

1 A related problem, the aliasing of parameters, occurs in languages that permit call by reference because
2 supposedly distinct parameters might refer to the same storage area, or a parameter and a non-local object might
3 refer to the same storage area. That vulnerability is described in Passing Parameters and Return Values [CSJ].

4 **6.48.4 Applicable language characteristics**

5 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 6 • A programming language that permits multiple interpretations of the same bit pattern.

7 **6.48.5 Avoiding the vulnerability or mitigating its effects**

8 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 9 • Programmers should avoid reinterpretation performed as a matter of convenience; for example, using an
10 integer pointer to manipulate character string data should be avoided. When type-breaking reinterpretation
11 is necessary, it should be carefully documented in the code. However this vulnerability cannot be
12 completely avoided because some applications view stored data in alternative ways.
- 13 • When using union types it is preferable to use discriminated unions. This is a form of a union where a
14 stored value indicates which interpretation is to be placed upon the data. Some languages (such as variant
15 records in Ada) enforce the view of data indicated by the value of the discriminant. If the language does not
16 enforce the interpretation (for example, equivalence in Fortran and union in C and C++), then the code
17 should implement an explicit discriminant and check its value before accessing the data in the union, or
18 use some other mechanism to ensure that correct interpretation is placed upon the data value.
- 19 • Operations that reinterpret the same stored value as representing a different type should be avoided. It is
20 easier to avoid such operations when the language clearly identifies them. For example, the name of Ada's
21 `Unchecked_Conversion` function explicitly warns of the problem. A much more difficult situation occurs
22 when pointers are used to achieve type reinterpretation. Some languages perform type-checking of
23 pointers and place restrictions on the ability of pointers to access arbitrary locations in storage. Others
24 permit the free use of pointers. In such cases, code must be carefully reviewed in a search for unintended
25 reinterpretation of stored values. Therefore it is important to explicitly comment the source code where
26 *intended* reinterpretations occur.
- 27 • Static analysis tools may be helpful in locating situations where unintended reinterpretation occurs. On the
28 other hand, the presence of reinterpretation greatly complicates static analysis for other problems, so it
29 may be appropriate to segregate intended reinterpretation operations into distinct subprograms.

30 **6.48.6 Implications for standardization**

31 In future standardization activities, the following items should be considered:

- 32 • Because the ability to perform reinterpretation is sometimes necessary, but the need for it is rare,
33 programming language designers might consider putting caution labels on operations that permit
34 reinterpretation. For example, the operation in Ada that permits unconstrained reinterpretation is called
35 `Unchecked_Conversion`.
- 36 • Because of the difficulties with undiscriminated unions, programming language designers might consider
37 offering union types that include distinct discriminants with appropriate enforcement of access to objects.
38

39 **6.49 Memory Leak [XYL]**

40 **6.49.1 Description of application vulnerability**

41 A memory leak occurs when software does not release allocated memory after it ceases to be used. Repeated
42 occurrences of a memory leak can consume considerable amounts of available memory. A memory leak can be
43 exploited by attackers to generate denial-of-service by causing the program to execute repeatedly a sequence that
44 triggers the leak. Moreover, a memory leak can cause any long-running critical program to shutdown prematurely.

1 6.49.2 Cross reference

2 CWE:

3 401. Failure to Release Memory Before Removing Last Reference (aka 'Memory Leak')

4 JSF AV Rule: 206

5 MISRA C 2004: 20.4

6 CERT C guidelines: MEM00-C and MEM31-C

7 Ada Quality and Style Guide: 5.4.5, 5.9.2, and 7.3.3

8 6.49.3 Mechanism of failure

9 As a process or system runs, any memory taken from dynamic memory and not returned or reclaimed (by the
10 runtime system or a garbage collector) after it ceases to be used, may result in future memory allocation requests
11 failing for lack of free space. Alternatively, memory claimed and returned can cause the heap to fragment, which
12 will eventually result in an inability to take the necessary size storage. Either condition will result in a memory
13 exhaustion exception, and program termination or a system crash.

14 If an attacker can determine the cause of an existing memory leak, the attacker may be able to cause the
15 application to leak quickly and therefore cause the application to crash.

16 6.49.4 Applicable language characteristics

17 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 18 • Languages that support mechanisms to dynamically allocate memory and reclaim memory under program
19 control.

20 6.49.5 Avoiding the vulnerability or mitigating its effects

21 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 22 • Use of garbage collectors that reclaim memory that will never be used by the application again. Some
23 garbage collectors are part of the language while others are add-ons
- 24 • Allocating and freeing memory in different modules and levels of abstraction may make it difficult for
25 developers to match requests to free storage with the appropriate storage allocation request. This may
26 cause confusion regarding when and if a block of memory has been allocated or freed, leading to memory
27 leaks. To avoid these situations, it is recommended that memory be allocated and freed at the same level
28 of abstraction, and ideally in the same code module.
- 29 • Storage pools are a specialized memory mechanism where all of the memory associated with a class of
30 objects is allocated from a specific bounded region. When used with strong typing one can ensure a strong
31 relationship between pointers and the space accessed such that storage exhaustion in one pool does not
32 affect the code operating on other memory.
- 33 • Memory leaks can be eliminated by avoiding the use of dynamically allocated storage entirely, or by doing
34 initial allocation exclusively and never allocating once the main execution commences. For safety-critical
35 systems and long running systems, the use of dynamic memory is almost always prohibited, or restricted to
36 the initialization phase of execution.
- 37 • Use static analysis that is capable of detecting when allocated storage is no longer used and has not been
38 freed (for reuse).

39 6.49.6 Implications for standardization

40 In future standardization activities, the following items should be considered:

- 41 • Languages can provide syntax and semantics to guarantee program-wide that dynamic memory is not
42 used (such as the configuration `pragmas` feature offered by some programming languages).

- Languages can document or can specify that implementations must document choices for dynamic memory management algorithms, to help designers decide on appropriate usage patterns and recovery techniques as necessary.

6.50 Argument Passing to Library Functions [TRJ]

6.50.1 Description of application vulnerability

Libraries that supply objects or functions are in most cases not required to check the validity of parameters passed to them. In those cases where parameter validation is required there might not be adequate parameter validation.

6.50.2 Cross reference

CWE:

114. Process Control

JSF AV Rules 16, 18, 19, 20, 21, 22, 23, 24, and 25

MISRA C 2004: 20.2, 20.3, 20.4, 20.6, 20.7, 20.8, 20.9, 20.10, 20.11, and 20.12

MISRA C++ 2008: 17-0-1, 17-0-5, 18-0-2, 18-0-3, 18-0-4, 18-2-1, 18-7-1 and 27-0-1

CERT C guidelines: INT03-C and STR07-C

6.50.3 Mechanism of failure

When calling a library, either the calling function or the library may make assumptions about parameters. For example, it may be assumed by a library that a parameter is non-zero so division by that parameter is performed without checking the value. Sometimes some validation is performed by the calling function, but the library may use the parameters in ways that were unanticipated by the calling function resulting in a potential vulnerability. Even when libraries do validate parameters, their response to an invalid parameter is usually undefined and can cause unanticipated results.

6.50.4 Applicable language characteristics

This vulnerability description is intended to be applicable to languages with the following characteristics:

- Languages providing or using libraries that do not validate the parameters accepted by functions, methods and objects.

6.50.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Libraries should be defined so that as many parameters as possible are validated.
- Libraries should be defined to validate any values passed to the library before the value is used.
- Develop wrappers around library functions that check the parameters before calling the function.
- Demonstrate statically that the parameters are never invalid.
- Use only libraries known to have been developed with consistent and validated interface requirements.

It is noted that several approaches can be taken, some work best if used in conjunction with each other.

6.50.6 Implications for standardization

In future standardization activities, the following items should be considered:

- All languages that define a support library should consider removing most if not all cases of undefined behaviour from the library sections.
- Libraries should be defined so that all parameters are validated.

1 **6.51 Dynamically-linked Code and Self-modifying Code [NYY]**

2 **6.51.1 Description of application vulnerability**

3 Code that is dynamically linked may be different from the code that was tested. This may be the result of replacing
4 a library with another of the same name or by altering an environment variable such as `LD_LIBRARY_PATH` on
5 UNIX platforms so that a different directory is searched for the library file. Executing code that is different than that
6 which was tested may lead to unanticipated errors or intentional malicious activity.

7 On some platforms, and in some languages, instructions can modify other instructions in the code space.
8 Historically self-modifying code was needed for software that was required to run on a platform with very limited
9 memory. It is now primarily used (or misused) to hide functionality of software and make it more difficult to reverse
10 engineer or for specialty applications such as graphics where the algorithm is tuned at runtime to give better
11 performance. Self-modifying code can be difficult to write correctly and even more difficult to test and maintain
12 correctly leading to unanticipated errors.

13 **6.51.2 Cross reference**

14 JSF AV Rule: 2

15 **6.51.3 Mechanism of failure**

16 Through the alteration of a library file or environment variable, the code that is dynamically linked may be different
17 from the code which was tested resulting in different functionality.

18 On some platforms, a pointer-to-data can erroneously be given an address value that designates a location in the
19 instruction space. If subsequently a modification is made through that pointer, then an unanticipated behaviour can
20 result.

21 **6.51.4 Applicable language characteristics**

22 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 23 • Languages that allow a pointer-to-data to be assigned an address value that designates a location in the
24 instruction space.
- 25 • Languages that allow execution of code that exists in data space.
- 26 • Languages that permit the use of dynamically linked or shared libraries.
- 27 • Languages that execute on an OS that permits program memory to be both writable and executable.

28 **6.51.5 Avoiding the vulnerability or mitigating its effects**

29 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 30 • Verify that the dynamically linked or shared code being used is the same as that which was tested.
- 31 • Do not write self-modifying code except in extremely rare instances. Most software applications should
32 never have a requirement for self-modifying code.
- 33 • In those extremely rare instances where its use is justified, self-modifying code should be very limited and
34 heavily documented.

35 **6.51.6 Implications for standardization**

36 In future standardization activities, the following items should be considered:

- 37 • Languages should consider providing a means so that a program can either automatically or manually
38 check that the digital signature of a library matches the one in the compile/test environment.

1 6.52 Library Signature [NSQ]

2 6.52.1 Description of application vulnerability

3 Programs written in modern languages may use libraries written in other languages than the program
4 implementation language. If the library is large, the effort of adding signatures for all of the functions use by hand
5 may be tedious and error-prone. Portable cross-language signatures will require detailed understanding of both
6 languages, which a programmer may lack.

7 Integrating two or more programming languages into a single executable relies upon knowing how to interface the
8 function calls, argument list and global data structures so the symbols match in the object code during linking.

9 Byte alignment can be a source of data corruption if memory boundaries between the programming languages are
10 different. Each language may also align structure data differently.

11 6.52.2 Cross reference

12 MISRA C 2004: 1.3

13 MISRA C++ 2008: 1-0-2

14 6.52.3 Mechanism of failure

15 When the library and the application in which it is to be used are written in different languages, the specification of
16 signatures is complicated by inter-language issues.

17 As used in this vulnerability description, the term library includes the interface to the operating system, which may
18 be specified only for the language used to code the operating system itself. In this case, any program written in
19 any other language faces the inter-language interoperability issue of creating a fully-functional signature.

20 When the application language and the library language are different, then the ability to specify signatures
21 according to either standard may not exist, or be very difficult. Thus, a translator-by-translator solution may be
22 needed, which maximizes the probability of incorrect signatures (since the solution must be recreated for each
23 translator pair). Incorrect signatures may or may not be caught during the linking phase.

24 6.52.4 Applicable language characteristics

25 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 26 • Languages that do not specify how to describe signatures for subprograms written in other languages.

27 6.52.5 Avoiding the vulnerability or mitigating its effects

28 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 29 • Use tools to create the signatures.
- 30 • Avoid using translator options or language features to reference library subprograms without proper
31 signatures.

32 6.52.6 Implications for standardization

33 In future standardization activities, the following items should be considered:

- 34 • Provide correct linkage even in the absence of correctly specified procedure signatures. (Note that this
35 may be very difficult where the original source code is unavailable.)
- 36 • Provide specified means to describe the signatures of subprograms.

1 **6.53 Unanticipated Exceptions from Library Routines [HJW]**

2 **6.53.1 Description of application vulnerability**

3 A library in this context is taken to mean a set of software routines produced outside the control of the main
4 application developer, usually by a third party, and where the application developer may not have access to the
5 source. In such circumstances the application developer has limited knowledge of the library functions, other than
6 from their behavioural interface.

7 Whilst the use of libraries can present a number of vulnerabilities, the focus of this vulnerability is any undesirable
8 behaviour that a library routine may exhibit, in particular the generation of unexpected exceptions.

9 **6.53.2 Cross reference**

10 JSF AV Rule: 208
11 MISRA C 2004: 3.6, 20.3
12 MISRA C++ 2008: 15-3-1, 15-3-2, 17-0-4
13 Ada Quaility and Style Guide: 5.8 and 7.5

14 **6.53.3 Mechanism of failure**

15 In some languages, unhandled exceptions lead to implementation-defined behaviour. This can include immediate
16 termination, without for example, releasing previously allocated resources. If a library routine raises an
17 unanticipated exception, this undesirable behaviour may result.

18 It should be noted that the considerations of [NZN], Returning Error Status, are also relevant here.

19 **6.53.4 Applicable language characteristics**

20 This vulnerability description is intended to be applicable to languages with the following characteristics:

- 21 • Languages that can link previously developed library code (where the developer and compiler don't have
22 access to the library source).
- 23 • Languages that permit exceptions to be thrown but do not require handlers for them.

24 **6.53.5 Avoiding the vulnerability or mitigating its effects**

25 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 26 • All library calls should be wrapped within a 'catch-all' exception handler (if the language supports such a
27 construct), so that any unanticipated exceptions can be caught and handled appropriately. This wrapping
28 may be done for each library function call or for the entire behaviour of the program, for example, having
29 the exception handler in main for C++. However, note that the latter isn't a complete solution, as static
30 objects are constructed before main is entered and are destroyed after it has been exited. Consequently,
31 MISRA C++ [16] bars class constructors and destructors from throwing exceptions (unless handled locally).
- 32 • An alternative approach would be to use only library routines for which all possible exceptions are
33 specified.

34 **6.53.6 Implications for standardization**

35 In future standardization activities, the following items should be considered:

- 36 • Languages that provide exceptions should provide a mechanism for catching all possible exceptions (for
37 example, a 'catch-all' handler). The behaviour of the program when encountering an unhandled exception
38 should be fully defined.

- 1 • Languages should provide a mechanism to determine which exceptions might be thrown by a called library
- 2 routine.

1 7. Application Vulnerabilities

2 7.1 Adherence to Least Privilege [XYN]

3 7.1.1 Description of application vulnerability

4 Failure to adhere to the principle of least privilege amplifies the risk posed by other vulnerabilities.

5 7.1.2 Cross reference

6 CWE:

7 250. Design Principle Violation: Failure to Use Least Privilege

8 CERT C guidelines: POS02-C

9 7.1.3 Mechanism of failure

10 This vulnerability type refers to cases in which an application grants greater access rights than necessary.
 11 Depending on the level of access granted, this may allow a user to access confidential information. For example,
 12 programs that run with root privileges have caused innumerable Unix security disasters. It is imperative that you
 13 carefully review privileged programs for all kinds of security problems, but it is equally important that privileged
 14 programs drop back to an unprivileged state as quickly as possible to limit the amount of damage that an
 15 overlooked vulnerability might be able to cause. Privilege management functions can behave in some less-than-
 16 obvious ways, and they have different quirks on different platforms. These inconsistencies are particularly
 17 pronounced if you are transitioning from one non-root user to another. Signal handlers and spawned processes run
 18 at the privilege of the owning process, so if a process is running as root when a signal fires or a sub-process is
 19 executed, the signal handler or sub-process will operate with root privileges. An attacker may be able to leverage
 20 these elevated privileges to do further damage. To grant the minimum access level necessary, first identify the
 21 different permissions that an application or user of that application will need to perform their actions, such as file
 22 read and write permissions, network socket permissions, and so forth. Then explicitly allow those actions while
 23 denying all else.

24 7.1.4 Avoiding the vulnerability or mitigating its effects

25 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 26 • Very carefully manage the setting, management and handling of privileges. Explicitly manage trust zones in
 27 the software.
- 28 • Follow the principle of least privilege when assigning access rights to entities in a software system.

29 7.2 Privilege Sandbox Issues [XYO]

30 7.2.1 Description of application vulnerability

31 A variety of vulnerabilities occur with improper handling, assignment, or management of privileges. These are
 32 especially present in sandbox environments, although it could be argued that any privilege problem occurs within
 33 the context of some sort of sandbox.

34 7.2.2 Cross reference

35 CWE:

36 266. Incorrect Privilege Assignment

37 267. Privilege Defined With Unsafe Actions

38 268. Privilege Chaining

39 269. Privilege Management Error

- 1 270. Privilege Context Switching Error
- 2 272. Least Privilege Violation
- 3 273. Failure to Check Whether Privileges were Dropped Successfully
- 4 274. Failure to Handle Insufficient Privileges
- 5 276. Insecure Default Permissions
- 6 CERT C guidelines: POS36-C

7 7.2.3 Mechanism of failure

8 The failure to drop system privileges when it is reasonable to do so is not an application vulnerability by itself. It
 9 does, however, serve to significantly increase the severity of other vulnerabilities. According to the principle of least
 10 privilege, access should be allowed only when it is absolutely necessary to the function of a given system, and only
 11 for the minimal necessary amount of time. Any further allowance of privilege widens the window of time during
 12 which a successful exploitation of the system will provide an attacker with that same privilege.

13 Many situations could lead to a mechanism of failure:

- 14 • A product could incorrectly assign a privilege to a particular entity.
- 15 • A particular privilege, role, capability, or right could be used to perform unsafe actions that were not
 16 intended, even when it is assigned to the correct entity. (Note that there are two separate sub-categories
 17 here: privilege incorrectly allows entities to perform certain actions; and the object is incorrectly accessible
 18 to entities with a given privilege.)
- 19 • Two distinct privileges, roles, capabilities, or rights could be combined in a way that allows an entity to
 20 perform unsafe actions that would not be allowed without that combination.
- 21 • The software may not properly manage privileges while it is switching between different contexts that cross
 22 privilege boundaries.
- 23 • A product may not properly track, modify, record, or reset privileges.
- 24 • In some contexts, a system executing with elevated permissions will hand off a process/file or other object
 25 to another process/user. If the privileges of an entity are not reduced, then elevated privileges are spread
 26 throughout a system and possibly to an attacker.
- 27 • The software may not properly handle the situation in which it has insufficient privileges to perform an
 28 operation.
- 29 • A program, upon installation, may set insecure permissions for an object.

30 7.2.4 Avoiding the vulnerability or mitigating its effects

31 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 32 • The principle of least privilege when assigning access rights to entities in a software system should be
 33 followed. The setting, management and handling of privileges should be managed very carefully. Upon
 34 changing security privileges, one should ensure that the change was successful.
- 35 • Consider following the principle of separation of privilege. Require multiple conditions to be met before
 36 permitting access to a system resource.
- 37 • Trust zones in the software should be explicitly managed. If at all possible, limit the allowance of system
 38 privilege to small, simple sections of code that may be called atomically.
- 39 • As soon as possible after acquiring elevated privilege to call a privileged function such as `chroot()`, the
 40 program should drop root privilege and return to the privilege level of the invoking user.
- 41 • In newer Windows implementations, make sure that the process token has the `SeImpersonatePrivilege`.

42 7.3 Executing or Loading Untrusted Code [XYS]

43 7.3.1 Description of application vulnerability

44 Executing commands or loading libraries from an untrusted source or in an untrusted environment can cause an
 45 application to execute malicious commands (and payloads) on behalf of an attacker.

1 7.3.2 Cross reference

2 CWE:

3 114. Process Control

4 CERT C guidelines: PRE09-C, ENV02-C, and ENV03-C

5 7.3.3 Mechanism of failure

6 Process control vulnerabilities take two forms:

- 7 • An attacker can change the command that the program executes so that the attacker explicitly controls
- 8 what the command is.
- 9 • An attacker can change the environment in which the command executes so that the attacker implicitly
- 10 controls what the command means.

11 Considering only the first scenario, the possibility that an attacker may be able to control the command that is

12 executed, process control vulnerabilities occur when:

- 13 • Data enters the application from a source that is not trusted.
- 14 • The data is used as or as part of a string representing a command that is executed by the application.
- 15 • By executing the command, the application gives an attacker a privilege or capability that the attacker
- 16 would not otherwise have.

17 7.3.4 Avoiding the vulnerability or mitigating its effects

18 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 19 • Libraries that are loaded should be well understood and come from a trusted source with a digital
- 20 signature. The application can execute code contained in native libraries, which often contain calls that
- 21 are susceptible to other security problems, such as buffer overflows or command injection.
- 22 • All native libraries should be validated.
- 23 • Determine if the application requires the use of the native library. It can be very difficult to determine what
- 24 these libraries actually do, and the potential for malicious code is high.
- 25 • To help prevent buffer overflow attacks, validate all input to native calls for content and length.
- 26 • If the native library does not come from a trusted source, review the source code of the library. The library
- 27 should be built from the reviewed source before using it.

28 7.4 Unspecified Functionality [BVQ]

29 7.4.1 Description of application vulnerability

30 *Unspecified functionality* is code that may be executed, but whose behaviour does not contribute to the

31 requirements of the application. While this may be no more than an amusing 'Easter Egg', like the flight simulator in

32 a spreadsheet, it does raise questions about the level of control of the development process.

33 In a security-critical environment particularly, the developer of an application could include a 'trap-door' to allow

34 illegitimate access to the system on which it is eventually executed, irrespective of whether the application has

35 obvious security requirements.

36 7.4.2 Cross reference

37 JSF AV Rule: 127

38 MISRA C 2004: 2.2, 2.3, 2.4, and 14.1

39 XYQ: Dead and Deactivated code.

1 7.4.3 Mechanism of failure

2 Unspecified functionality is not a software vulnerability per se, but more a development issue. In some cases,
3 unspecified functionality may be added by a developer without the knowledge of the development organization. In
4 other cases, typically Easter Eggs, the functionality is unspecified as far as the user is concerned (nobody buys a
5 spreadsheet expecting to find it includes a flight simulator), but is specified by the development organization. In
6 effect they only reveal a subset of the program's behaviour to the users.

7 In the first case, one would expect a well managed development environment to discover the additional
8 functionality during validation and verification. In the second case, the user is relying on the supplier not to release
9 harmful code.

10 In effect, a program's requirements are 'the program should behave in the following manner and do nothing
11 else'. The 'and do nothing else' clause is often not explicitly stated, and can be difficult to demonstrate.

12 7.4.4 Avoiding the vulnerability or mitigating its effects

13 End users can avoid the vulnerability or mitigate its ill effects in the following ways:

- 14 • Programs and development tools that are to be used in critical applications should come from a developer
15 who uses a recognized and audited development process for the development of those programs and
16 tools. For example: ISO 9001 or CMMI®.
- 17 • The development process should generate documentation showing traceability from source code to
18 requirements, in effect answering 'why is this unit of code in this program?'. Where unspecified functionality
19 is there for a legitimate reason (such as diagnostics required for developer maintenance or enhancement),
20 the documentation should also record this. It is not unreasonable for customers of bespoke critical code to
21 ask to see such traceability as part of their acceptance of the application

22 7.5 Distinguished Values in Data Types [KLK]

23 7.5.1 Description of application vulnerability

24 Sometimes, in a type representation, certain values are distinguished as not being members of the type, but rather
25 as providing auxiliary information. Examples include special characters used as string terminators, distinguished
26 values used to indicate out of type entries in *SQL* (Structured Query Language) database fields, and sentinels used
27 to indicate the bounds of queues or other data structures. When the usage pattern of code containing distinguished
28 values is changed, it may happen that the distinguished value happens to coincide with a legitimate in-type value.
29 In such a case, the value is no longer distinguishable from an in-type value and the software will no longer produce
30 the intended results.

31 7.5.2 Cross reference

32 CWE
33 20: Improper input validation
34 137: Representation errors
35 JSF AV Rule: 151

36 7.5.3 Mechanism of failure

37 A "distinguished value" or a "magic number" in the representation of a data type might be used to represent out-of-
38 type information. Some examples include the following:

- 39 • The use of a special code, such as "00", to indicate the termination of a coded character string.
- 40 • The use of a special value, such as "999...9", as the indication that the actual value is either not known or
41 is invalid.

1 If the use of the software is later generalized, the once-special value can become indistinguishable from valid data.
 2 Note that the problem may occur simply if the pattern of usage of the software is changed from that anticipated by
 3 the software's designers. It may also occur if the software is reused in other circumstances.

4 An example of a change in the pattern of usage is this: An organization logs visitors to its buildings by recording
 5 their names and national identity numbers or social security numbers in a database. Of course, some visitors
 6 legitimately don't have or don't know their social security number, so the receptionists enter numbers to "make the
 7 computer happy." Receptionists at one building have adopted the convention of using the code "555-55-5555" to
 8 designate children of employees. Receptionists at another building have used the same code to designate foreign
 9 nationals. When the databases are merged, the children are reclassified as foreign nationals or vice-versa
 10 depending on which set of receptionists are using the newly merged database.

11 An example of an unanticipated change due to reuse is this: Suppose a software component analyzes radar data,
 12 recording data every degree of azimuth from 0 to 359. Packets of data are sent to other components for
 13 processing, updating displays, recording, and so on. Since all degree values are non-negative, a distinguished
 14 value of -1 is used as a signal to stop processing, compute summary data, close files, and so on. Many of the
 15 components are to be reused in a new system with a new radar analysis component. However the new component
 16 represents direction by numbers in the range -180 degrees to 179 degrees. When an azimuth value of -1 is
 17 provided, the downstream components will interpret that as the indication to stop processing. If the magic value is
 18 changed to, say, -999, the software is still at risk of failing when future enhancements (say, counting accumulated
 19 degrees on complete revolutions) bring -999 into the range of valid data.

20 Distinguished values should be avoided. Instead, the software should be designed to use distinct variables to
 21 encode the desired out-of-type information. For example, the length of a character string might be encoded in a
 22 dope vector and validity of data entries might be encoded in distinct Boolean values.

23 **7.5.4 Avoiding the vulnerability or mitigating its effects**

24 End users can avoid the vulnerability or mitigate its ill effects in the following ways:

- 25 • Use auxiliary variables (perhaps enclosed in variant records) to encode out-of-type information.
- 26 • Use enumeration types to convey category information. Do not rely upon large ranges of integers, with
 27 distinguished values having special meanings.
- 28 • Use named constants to make it easier to change distinguished values.

29 **7.6 Memory Locking [XZX]**

30 **7.6.1 Description of application vulnerability**

31 Sensitive data stored in memory that was not locked or that has been improperly locked may be written to swap
 32 files on disk by the virtual memory manager.

33 **7.6.2 Cross reference**

34 CWE:
 35 591. Sensitive Data Storage in Improperly Locked Memory
 36 CERT C guidelines: MEM06-C

37 **7.6.3 Mechanism of failure**

38 Sensitive data that is not kept cryptographically secure may become visible to an attacker by any of several
 39 mechanisms. Some operating systems may write memory to swap or page files that may be visible to an attacker.
 40 Some operating systems may provide mechanisms to examine the physical memory of the system or the virtual
 41 memory of another application. Application debuggers may be able to stop the target application and examine or
 42 alter memory.

7.6.4 Avoiding the vulnerability or mitigating its effects

In almost all cases, these attacks require elevated or appropriate privilege.

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Remove debugging tools from production systems.
- Log and audit all privileged operations.
- Identify data that needs to be protected and use appropriate cryptographic and other data obfuscation techniques to avoid keeping plaintext versions of this data in memory or on disk.
- If the operating system allows, clear the swap file on shutdown.

Note: Several implementations of the POSIX `mlock()` and the Microsoft Windows `VirtualLock()` functions will prevent the named memory region from being written to a swap or page file. However, such usage is not portable.

Systems that provide a "hibernate" facility (such as laptops) will write all of physical memory to a file that may be visible to an attacker on resume.

7.7 Resource Exhaustion [XZP]

7.7.1 Description of application vulnerability

The application is susceptible to generating and/or accepting an excessive number of requests that could potentially exhaust limited resources, such as memory, file system storage, database connection pool entries, or CPU. This could ultimately lead to a denial of service that could prevent any other applications from accessing these resources.

7.7.2 Cross reference

CWE:
400. Resource Exhaustion

7.7.3 Mechanism of failure

There are two primary failures associated with resource exhaustion. The most common result of resource exhaustion is denial of service. In some cases an attacker or a defect may cause a system to fail in an unsafe or insecure fashion by causing an application to exhaust the available resources.

Resource exhaustion issues are generally understood but are far more difficult to prevent. Taking advantage of various entry points, an attacker could craft a wide variety of requests that would cause the site to consume resources. Database queries that take a long time to process are good *DoS* (Denial of Service) targets. An attacker would only have to write a few lines of Perl code to generate enough traffic to exceed the site's ability to keep up. This would effectively prevent authorized users from using the site at all.

Resources can be exhausted simply by ensuring that the target machine must do much more work and consume more resources to service a request than the attacker must do to initiate a request. Prevention of these attacks requires either that the target system either recognizes the attack and denies that user further access for a given amount of time or uniformly throttles all requests to make it more difficult to consume resources more quickly than they can again be freed. The first of these solutions is an issue in itself though, since it may allow attackers to prevent the use of the system by a particular valid user. If the attacker impersonates the valid user, he may be able to prevent the user from accessing the server in question. The second solution is simply difficult to effectively institute and even when properly done, it does not provide a full solution. It simply makes the attack require more resources on the part of the attacker.

The final concern that must be discussed about issues of resource exhaustion is that of systems which "fail open." This means that in the event of resource consumption, the system fails in such a way that the state of the system

1 — and possibly the security functionality of the system — are compromised. A prime example of this can be found
 2 in old switches that were vulnerable to "macof" attacks (so named for a tool developed by Dugsong). These attacks
 3 flooded a switch with random IP and MAC address combinations, therefore exhausting the switch's cache, which
 4 held the information of which port corresponded to which MAC addresses. Once this cache was exhausted, the
 5 switch would fail in an insecure way and would begin to act simply as a hub, broadcasting all traffic on all ports and
 6 allowing for basic sniffing attacks.

7 7.7.4 Avoiding the vulnerability or mitigating its effects

8 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 9 • Implement throttling mechanisms into the system architecture. The best protection is to limit the amount of
 10 resources that an application can cause to be expended. A strong authentication and access control model
 11 will help prevent such attacks from occurring in the first place. The authentication application should be
 12 protected against denial of service attacks as much as possible. Limiting the database access, perhaps by
 13 caching result sets, can help minimize the resources expended. To further limit the potential for a denial of
 14 service attack, consider tracking the rate of requests received from users and blocking requests that
 15 exceed a defined rate threshold.
- 16 • Ensure that applications have specific limits of scale placed on them, and ensure that all failures in
 17 resource allocation cause the application to fail safely.

18 7.8 Injection [RST]

19 7.8.1 Description of application vulnerability

20 Injection problems span a wide range of instantiations. The basic form of this weakness involves the software
 21 allowing injection of additional data in input data to alter the control flow of the process. Command injection
 22 problems are a subset of injection problem, in which the process can be tricked into calling external processes of
 23 an attacker's choice through the injection of command syntax into the input data. Multiple leading/internal/trailing
 24 special elements injected into an application through input can be used to compromise a system. As data is parsed,
 25 improperly handled multiple leading special elements may cause the process to take unexpected actions that result
 26 in an attack. Software may allow the injection of special elements that are non-typical but equivalent to typical
 27 special elements with control implications. This frequently occurs when the product has protected itself against
 28 special element injection. Software may allow inputs to be fed directly into an output file that is later processed as
 29 code, such as a library file or template. Line or section delimiters injected into an application can be used to
 30 compromise a system.

31 Many injection attacks involve the disclosure of important information -- in terms of both data sensitivity and
 32 usefulness in further exploitation. In some cases injectable code controls authentication; this may lead to a remote
 33 vulnerability. Injection attacks are characterized by the ability to significantly change the flow of a given process,
 34 and in some cases, to the execution of arbitrary code. Data injection attacks lead to loss of data integrity in nearly
 35 all cases as the control-plane data injected is always incidental to data recall or writing. Often the actions performed
 36 by injected control code are not logged.

37 SQL injection attacks are a common instantiation of injection attack, in which SQL commands are injected into
 38 input to effect the execution of predefined SQL commands. Since SQL databases generally hold sensitive data,
 39 loss of confidentiality is a frequent problem with SQL injection vulnerabilities. If poorly implemented SQL
 40 commands are used to check user names and passwords, it may be possible to connect to a system as another
 41 user with no previous knowledge of the password. If authorization information is held in a SQL database, it may be
 42 possible to change this information through the successful exploitation of the SQL injection vulnerability. Just as it
 43 may be possible to read sensitive information, it is also possible to make changes or even delete this information
 44 with a SQL injection attack.

45 Injection problems encompass a wide variety of issues — all mitigated in very different ways. The most important
 46 issue to note is that all injection problems share one thing in common — they allow for the injection of control data
 47 into the user controlled data. This means that the execution of the process may be altered by sending code in
 48 through legitimate data channels, using no other mechanism. While buffer overflows and many other flaws involve

1 the use of some further issue to gain execution, injection problems need only for the data to be parsed. Many
 2 injection attacks involve the disclosure of important information in terms of both data sensitivity and usefulness in
 3 further exploitation. In some cases injectable code controls authentication, this may lead to a remote vulnerability.

4 **7.8.2 Cross reference**

5 CWE:

- 6 74. Failure to Sanitize Data into a Different Plane ('Injection')
- 7 76. Failure to Resolve Equivalent Special Elements into a Different Plane
- 8 78. Failure to Sanitize Data into an OS Command (aka 'OS Command Injection')
- 9 90. Failure to Sanitize Data into LDAP Queries (aka 'LDAP Injection')
- 10 91. XML Injection (aka Blind XPath Injection)
- 11 92. Custom Special Character Injection
- 12 95. Insufficient Control of Directives in Dynamically Code Evaluated Code (aka 'Eval Injection')
- 13 97. Failure to Sanitize Server-Side Includes (SSI) Within a Web Page
- 14 98. Insufficient Control of Filename for Include/Require Statement in PHP Program (aka 'PHP File Inclusion')
- 15 99. Insufficient Control of Resource Identifiers (aka 'Resource Injection')
- 16 144. Failure to Sanitize Line Delimiters
- 17 145. Failure to Sanitize Section Delimiters
- 18 161. Failure to Sanitize Multiple Leading Special Elements
- 19 163. Failure to Sanitize Multiple Trailing Special Elements
- 20 165. Failure to Sanitize Multiple Internal Special Elements
- 21 166. Failure to Handle Missing Special Element
- 22 167. Failure to Handle Additional Special Element
- 23 168. Failure to Resolve Inconsistent Special Elements
- 24 564. SQL Injection: Hibernate

25 CERT C guidelines: FIO30-C

26 **7.8.3 Mechanism of failure**

27 A software system that accepts and executes input in the form of operating system commands (such as
 28 `system()`, `exec()`, `open()`) could allow an attacker with lesser privileges than the target software to execute
 29 commands with the elevated privileges of the executing process. Command injection is a common problem with
 30 wrapper programs. Often, parts of the command to be run are controllable by the end user. If a malicious user
 31 injects a character (such as a semi-colon) that delimits the end of one command and the beginning of another, he
 32 may then be able to insert an entirely new and unrelated command to do whatever he pleases.

33 Dynamically generating operating system commands that include user input as parameters can lead to command
 34 injection attacks. An attacker can insert operating system commands or modifiers in the user input that can cause
 35 the request to behave in an unsafe manner. Such vulnerabilities can be very dangerous and lead to data and
 36 system compromise. If no validation of the parameter to the `exec` command exists, an attacker can execute any
 37 command on the system the application has the privilege to access.

38 There are two forms of command injection vulnerabilities. An attacker can change the command that the program
 39 executes (the attacker explicitly controls what the command is). Alternatively, an attacker can change the
 40 environment in which the command executes (the attacker implicitly controls what the command means). The first
 41 scenario where an attacker explicitly controls the command that is executed can occur when:

- 42 • Data enters the application from an untrusted source.
- 43 • The data is part of a string that is executed as a command by the application.
- 44 • By executing the command, the application gives an attacker a privilege or capability that the attacker
 45 would not otherwise have.

46 Eval injection occurs when the software allows inputs to be fed directly into a function (such as "eval") that
 47 dynamically evaluates and executes the input as code, usually in the same interpreted language that the product
 48 uses. Eval injection is prevalent in handler/dispatch procedures that might want to invoke a large number of
 49 functions, or set a large number of variables.

1 A PHP file inclusion occurs when a PHP product uses `require` or `include` statements, or equivalent statements,
 2 that use attacker-controlled data to identify code or *HTML* (HyperText Markup Language) to be directly processed
 3 by the PHP interpreter before inclusion in the script.

4 A resource injection issue occurs when the following two conditions are met:

- 5 • An attacker can specify the identifier used to access a system resource. For example, an attacker might be
 6 able to specify part of the name of a file to be opened or a port number to be used.
- 7 • By specifying the resource, the attacker gains a capability that would not otherwise be permitted. For
 8 example, the program may give the attacker the ability to overwrite the specified file, run with a
 9 configuration controlled by the attacker, or transmit sensitive information to a third-party server. Note:
 10 Resource injection that involves resources stored on the file system goes by the name path manipulation
 11 and is reported in separate category. See Path Traversal [EWR] description for further details of this
 12 vulnerability. Allowing user input to control resource identifiers may enable an attacker to access or modify
 13 otherwise protected system resources.

14 Line or section delimiters injected into an application can be used to compromise a system. As data is parsed, an
 15 injected/absent/malformed delimiter may cause the process to take unexpected actions that result in an attack. One
 16 example of a section delimiter is the boundary string in a multipart MIME message. In many cases, doubled line
 17 delimiters can serve as a section delimiter.

18 **7.8.4 Avoiding the vulnerability or mitigating its effects**

19 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 20 • Assume all input is malicious. Use an appropriate combination of black lists and white lists to ensure only
 21 valid, expected and appropriate input is processed by the system.
- 22 • Narrowly define the set of safe characters based on the expected values of the parameter in the request.
- 23 • Developers should anticipate that delimiters and special elements would be injected/removed/manipulated
 24 in the input vectors of their software system and appropriate mechanisms should be put in place to handle
 25 them.
- 26 • Implement SQL strings using prepared statements that bind variables. Prepared statements that do not
 27 bind variables can be vulnerable to attack.
- 28 • Use vigorous white-list style checking on any user input that may be used in a SQL command. Rather than
 29 escape meta-characters, it is safest to disallow them entirely since the later use of data that have been
 30 entered in the database may neglect to escape meta-characters before use.
- 31 • Follow the principle of least privilege when creating user accounts to a SQL database. Users should only
 32 have the minimum privileges necessary to use their account. If the requirements of the system indicate that
 33 a user can read and modify their own data, then limit their privileges so they cannot read/write others' data.
- 34 • Assign permissions to the software system that prevents the user from accessing/opening privileged files.
- 35 • To avert eval injections, refactor your code so that it does not need to use `eval()`.

36 **7.9 Cross-site Scripting [XYT]**

37 **7.9.1 Description of application vulnerability**

38 *Cross-site scripting* (XSS) occurs when dynamically generated web pages display input, such as login information
 39 that is not properly validated, allowing an attacker to embed malicious scripts into the generated page and then
 40 execute the script on the machine of any user that views the site. If successful, cross-site scripting vulnerabilities
 41 can be exploited to manipulate or steal cookies, create requests that can be mistaken for those of a valid user,
 42 compromise confidential information, or execute malicious code on the end user systems for a variety of nefarious
 43 purposes.

44 **7.9.2 Cross reference**

45 CWE:

- 1 79. Failure to Preserve Web Page Structure ('Cross-site Scripting')
- 2 80. Failure to Sanitize Script-Related HTML Tags in a Web Page (Basic XSS)
- 3 81. Failure to Sanitize Directives in an Error Message Web Page
- 4 82. Failure to Sanitize Script in Attributes of IMG Tags in a Web Page
- 5 83. Failure to Sanitize Script in Attributes in a Web Page
- 6 84. Failure to Resolve Encoded URI Schemes in a Web Page
- 7 85. Doubled Character XSS Manipulations
- 8 86. Invalid Characters in Identifiers
- 9 87. Alternate XSS Syntax

10 7.9.3 Mechanism of failure

11 Cross-site scripting (XSS) vulnerabilities occur when an attacker uses a web application to send malicious code,
 12 generally JavaScript, to a different end user. When a web application uses input from a user in the output it
 13 generates without filtering it, an attacker can insert an attack in that input and the web application sends the attack
 14 to other users. The end user trusts the web application, and the attacks exploit that trust to do things that would not
 15 normally be allowed. Attackers frequently use a variety of methods to encode the malicious portion of the tag, such
 16 as using Unicode, so the request looks less suspicious to the user.

17 XSS attacks can generally be categorized into two categories: stored and reflected. Stored attacks are those where
 18 the injected code is permanently stored on the target servers in a database, message forum, visitor log, and so
 19 forth. Reflected attacks are those where the injected code takes another route to the victim, such as in an email
 20 message, or on some other server. When a user is tricked into clicking a link or submitting a form, the injected code
 21 travels to the vulnerable web server, which reflects the attack back to the user's browser. The browser then
 22 executes the code because it came from a 'trusted' server. For a reflected XSS attack to work, the victim must
 23 submit the attack to the server. This is still a very dangerous attack given the number of possible ways to trick a
 24 victim into submitting such a malicious request, including clicking a link on a malicious Web site, in an email, or in
 25 an inter-office posting.

26 XSS flaws are very common in web applications, as they require a great deal of developer discipline to avoid them
 27 in most applications. It is relatively easy for an attacker to find XSS vulnerabilities. Some of these vulnerabilities can
 28 be found using scanners, and some exist in older web application servers. The consequence of an XSS attack is
 29 the same regardless of whether it is stored or reflected.

30 The difference is in how the payload arrives at the server. XSS can cause a variety of problems for the end user
 31 that range in severity from an annoyance to complete account compromise. The most severe XSS attacks involve
 32 disclosure of the user's session cookie, which allows an attacker to hijack the user's session and take over their
 33 account. Other damaging attacks include the disclosure of end user files, installation of Trojan horse programs,
 34 redirecting the user to some other page or site, and modifying presentation of content.

35 Cross-site scripting (XSS) vulnerabilities occur when:

- 36 • Data enters a Web application through an untrusted source, most frequently a web request. The data is
 37 included in dynamic content that is sent to a web user without being validated for malicious code.
- 38 • The malicious content sent to the web browser often takes the form of a segment of JavaScript, but may
 39 also include HTML, Flash or any other type of code that the browser may execute. The variety of attacks
 40 based on XSS is almost limitless, but they commonly include transmitting private data like cookies or other
 41 session information to the attacker, redirecting the victim to web content controlled by the attacker, or
 42 performing other malicious operations on the user's machine under the guise of the vulnerable site.

43 Cross-site scripting attacks can occur wherever an untrusted user has the ability to publish content to a trusted web
 44 site. Typically, a malicious user will craft a client-side script, which — when parsed by a web browser — performs
 45 some activity (such as sending all site cookies to a given e-mail address). If the input is unchecked, this script will
 46 be loaded and run by each user visiting the web site. Since the site requesting to run the script has access to the
 47 cookies in question, the malicious script does also. There are several other possible attacks, such as running
 48 "Active X" controls (under Microsoft Internet Explorer) from sites that a user perceives as trustworthy; cookie theft is
 49 however by far the most common. All of these attacks are easily prevented by ensuring that no script tags — or for
 50 good measure, HTML tags at all — are allowed in data to be posted publicly.

1 Specific instances of XSS are:

- 2 • 'Basic' XSS involves a complete lack of cleansing of any special characters, including the most
- 3 fundamental XSS elements such as "<", ">", and "&".
- 4 • A web developer displays input on an error page (such as a customized 403 Forbidden page). If an
- 5 attacker can influence a victim to view/request a web page that causes an error, then the attack may be
- 6 successful.
- 7 • A Web application that trusts input in the form of HTML IMG tags is potentially vulnerable to XSS attacks.
- 8 Attackers can embed XSS exploits into the values for IMG attributes (such as SRC) that is streamed and
- 9 then executed in a victim's browser. Note that when the page is loaded into a user's browser, the exploit
- 10 will automatically execute.
- 11 • The software does not filter "javascript:" or other *URI*'s (Uniform Resource Identifier) from dangerous
- 12 attributes within tags, such as *onmouseover*, *onload*, *onerror*, or *style*.
- 13 • The web application fails to filter input for executable script disguised with URI encodings.
- 14 • The web application fails to filter input for executable script disguised using doubling of the involved
- 15 characters.
- 16 • The software does not strip out invalid characters in the middle of tag names, schemes, and other
- 17 identifiers, which are still rendered by some web browsers that ignore the characters.
- 18 • The software fails to filter alternate script syntax provided by the attacker.

19 Cross-site scripting attacks may occur anywhere that possibly malicious users are allowed to post unregulated
 20 material to a trusted web site for the consumption of other valid users. The most common example can be found in
 21 bulletin-board web sites that provide web based mailing list-style functionality. The most common attack performed
 22 with cross-site scripting involves the disclosure of information stored in user cookies. In some circumstances it may
 23 be possible to run arbitrary code on a victim's computer when cross-site scripting is combined with other flaws.

24 **7.9.4 Avoiding the vulnerability or mitigating its effects**

25 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 26 • Carefully check each input parameter against a rigorous positive specification (white list) defining the
- 27 specific characters and format allowed.
- 28 • All input should be sanitized, not just parameters that the user is supposed to specify, but all data in the
- 29 request, including hidden fields, cookies, headers, the *URL* (Uniform Resource Locator) itself, and so forth.
- 30 • A common mistake that leads to continuing XSS vulnerabilities is to validate only fields that are expected
- 31 to be redisplayed by the site.
- 32 • Data is frequently encountered from the request that is reflected by the application server or the
- 33 application that the development team did not anticipate. Also, a field that is not currently reflected may be
- 34 used by a future developer. Therefore, validating ALL parts of the HTTP request is recommended.

35 **7.10 Unquoted Search Path or Element [XZQ]**

36 **7.10.1 Description of application vulnerability**

37 Strings injected into a software system that are not quoted can permit an attacker to execute arbitrary commands.

38 **7.10.2 Cross reference**

39 CWE:

40 428. Unquoted Search Path or Element

41 CERT C guidelines: ENV04-C

42 **7.10.3 Mechanism of failure**

43 The mechanism of failure stems from missing quoting of strings injected into a software system. By allowing white-
 44 spaces in identifiers, an attacker could potentially execute arbitrary commands. This vulnerability covers

"C:\Program Files" and space-in-search-path issues. Theoretically this could apply to other operating systems besides Windows, especially those that make it easy for spaces to be in filenames or folders names.

7.10.4 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Software should quote the input data that can be potentially executed on a system.
- Use a programming language that enforces the quoting of strings.

7.11 Improperly Verified Signature [XZR]

7.11.1 Description of application vulnerability

The software does not verify, or improperly verifies, the cryptographic signature for data. By not adequately performing the verification step, the data being received should not be trusted and may be corrupted or made intentionally incorrect by an adversary.

7.11.2 Cross reference

CWE:
347. Improperly Verified Signature

7.11.3 Mechanism of failure

Data is signed using techniques that assure the integrity of the data. There are two ways that the integrity can be intentionally compromised. The exchange of the cryptologic keys may have been compromised so that an attacker could provide encrypted data that has been altered. Alternatively, the cryptologic verification could be flawed so that the encryption of the data is flawed which again allows an attacker to alter the data.

7.11.4 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Use data signatures to the extent possible to help ensure trust in data.
- Use built-in verifications for data.

7.12 Discrepancy Information Leak [XZL]

7.12.1 Description of application vulnerability

A discrepancy information leak is an information leak in which the product behaves differently, or sends different responses, in a way that reveals security-relevant information about the state of the product, such as whether a particular operation was successful or not.

7.12.2 Cross reference

CWE:
203. Discrepancy Information Leaks
204. Response Discrepancy Information Leak
206. Internal Behavioural Inconsistency Information Leak
207. External Behavioral Inconsistency Information Leak
208. Timing Discrepancy Information Leak

1 **7.12.3 Mechanism of failure**

2 A response discrepancy information leak occurs when the product sends different messages in direct response to
3 an attacker's request, in a way that allows the attacker to learn about the inner state of the product. The leaks can
4 be inadvertent (bug) or intentional (design).

5
6 A behavioural discrepancy information leak occurs when the product's actions indicate important differences based
7 on (1) the internal state of the product or (2) differences from other products in the same class. Attacks such as OS
8 fingerprinting rely heavily on both behavioural and response discrepancies. An internal behavioural inconsistency
9 information leak is the situation where two separate operations in a product cause the product to behave differently
10 in a way that is observable to an attacker and reveals security-relevant information about the internal state of the
11 product, such as whether a particular operation was successful or not. An external behavioural inconsistency
12 information leak is the situation where the software behaves differently than other products like it, in a way that is
13 observable to an attacker and reveals security-relevant information about which product is being used, or its
14 operating state.

15
16 A timing discrepancy information leak occurs when two separate operations in a product require different amounts
17 of time to complete, in a way that is observable to an attacker and reveals security-relevant information about the
18 state of the product, such as whether a particular operation was successful or not.

19 **7.12.4 Avoiding the vulnerability or mitigating its effects**

20 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 21 • Compartmentalize the system to have "safe" areas where trust boundaries can be unambiguously drawn.
- 22 • Do not allow sensitive data to go outside of the trust boundary and always be careful when interfacing with
- 23 a compartment outside of the safe area.

24 **7.13 Sensitive Information Uncleared Before Release [XZK]**

25 **7.13.1 Description of application vulnerability**

26 The software does not fully clear previously used information in a data structure, file, or other resource, before
27 making that resource available to another party that did not have access to the original information.

28 **7.13.2 Cross reference**

29 CWE:

30 226. Sensitive Information Uncleared Before Release

31 CERT C guidelines: MEM03-C

32 **7.13.3 Mechanism of failure**

33 This typically involves memory in which the new data occupies less memory than the old data, which leaves
34 portions of the old data still available ("memory disclosure"). However, equivalent errors can occur in other
35 situations where the length of data is variable but the associated data structure is not. This can overlap with
36 cryptographic errors and cross-boundary cleansing info leaks.

37 Dynamic memory managers are not required to clear freed memory and generally do not because of the additional
38 runtime overhead. Furthermore, dynamic memory managers are free to reallocate this same memory. As a result,
39 it is possible to accidentally leak sensitive information if it is not cleared before calling a function that frees dynamic
40 memory. Programmers should not and can not rely on memory being cleared during allocation.

41 **7.13.4 Avoiding the vulnerability or mitigating its effects**

42 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Use library functions and or programming language features (such as destructors or finalization procedures) that provide automatic clearing of freed buffers and or the functionality to clear buffers.

7.14 Path Traversal [EWR]

7.14.1 Description of application vulnerability

The software constructs a path that contains relative traversal sequence such as ".." or an absolute path sequence such as "/path/here." Attackers run the software in a particular directory so that the hard link or symbolic link used by the software accesses a file that the attacker has under their control. In doing this, the attacker may be able to escalate their privilege level to that of the running process.

7.14.2 Cross reference

CWE:

- 22. Path Traversal
 - 24. Path Traversal: - '../filedir'
 - 25. Path Traversal: '/../filedir'
 - 26. Path Traversal: '/dir../filename'
 - 27. Path Traversal: 'dir../filename'
 - 28. Path Traversal: '..filename'
 - 29. Path Traversal: '\\.filename'
 - 30. Path Traversal: 'dir\\.filename'
 - 31. Path Traversal: 'dir\\.filename'
 - 32. Path Traversal: '...' (Triple Dot)
 - 33. Path Traversal: '....' (Multiple Dot)
 - 34. Path Traversal: '.../'
 - 35. Path Traversal: '.../.../'
 - 37. Path Traversal: '/absolute/pathname/here'
 - 38. Path Traversal: '\\absolute\\pathname\\here'
 - 39. Path Traversal: 'C:dirname'
 - 40. Path Traversal: '\\UNC\\share\\name\\' (Windows UNC Share)
 - 61. UNIX Symbolic Link (Symlink) Following
 - 62. UNIX Hard Link
 - 64. Windows Shortcut Following (.LNK)
 - 65. Windows Hard Link
- CERT C guidelines: FIO02-C

7.14.3 Mechanism of failure

There are two primary ways that an attacker can orchestrate an attack using path traversal. In the first, the attacker alters the path being used by the software to point to a location that the attacker has control over. Alternatively, the attacker has no control over the path, but can alter the directory structure so that the path points to a location that the attacker does have control over.

For instance, a software system that accepts input in the form of: '..filename', '\\.filename', '/directory../filename', 'directory../filename', '..filename', '\\.filename', 'directory\\.filename', 'directory\\.\\.filename', '...', '....' (multiple dots), '.../' or '.../.../' without appropriate validation can allow an attacker to traverse the file system to access an arbitrary file. Note that '..' is ignored if the current working directory is the root directory. Some of these input forms can be used to cause problems for systems that strip out '..' from input in an attempt to remove relative path traversal.

There are several common ways that an attacker can point a file access to a file the attacker has under their control. A software system that accepts input in the form of '/absolute/pathname/here' or '\\absolute\\pathname\\here' without appropriate validation can also allow an attacker to traverse the file system to unintended locations or access arbitrary files. An attacker can inject a drive letter or Windows volume letter ('C:dirname') into a software system to potentially redirect access to an unintended location or arbitrary file. A software system that accepts input

1 in the form of a backslash absolute path without appropriate validation can allow an attacker to traverse the file
 2 system to unintended locations or access arbitrary files. An attacker can inject a Windows UNC share
 3 ("\\UNC\share\name") into a software system to potentially redirect access to an unintended location or arbitrary file.
 4 A software system that allows UNIX symbolic links (symlink) as part of paths whether in internal code or through
 5 user input can allow an attacker to spoof the symbolic link and traverse the file system to unintended locations or
 6 access arbitrary files. The symbolic link can permit an attacker to read/write/corrupt a file that they originally did not
 7 have permissions to access. Failure for a system to check for hard links can result in vulnerability to different types
 8 of attacks. For example, an attacker can escalate their privileges if he/she can replace a file used by a privileged
 9 program with a hard link to a sensitive file, for example, `etc/passwd`. When the process opens the file, the
 10 attacker can assume the privileges of that process.

11 A software system that allows Windows shortcuts (.LNK) as part of paths whether in internal code or through user
 12 input can allow an attacker to spoof the symbolic link and traverse the file system to unintended locations or access
 13 arbitrary files. The shortcut (file with the .lnk extension) can permit an attacker to read/write a file that they originally
 14 did not have permissions to access.

15 Failure for a system to check for hard links can result in vulnerability to different types of attacks. For example, an
 16 attacker can escalate their privileges if he/she can replace a file used by a privileged program with a hard link to a
 17 sensitive file (such as `etc/passwd`). When the process opens the file, the attacker can assume the privileges of
 18 that process or possibly prevent a program from accurately processing data in a software system.

19 7.14.4 Avoiding the vulnerability or mitigating its effects

20 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 21 • Assume all input is malicious. Attackers can insert paths into input vectors and traverse the file system.
- 22 • Use an appropriate combination of black lists and white lists to ensure only valid and expected input is
 23 processed by the system.
- 24 • A sanitizing mechanism can remove characters such as '.' and ';' which may be required for some exploits.
 25 An attacker can try to fool the sanitizing mechanism into "cleaning" data into a dangerous form. Suppose
 26 the attacker injects a '.' inside a filename (say, "sensitiveFile") and the sanitizing mechanism removes the
 27 character resulting in the valid filename, "sensitiveFile". If the input data are now assumed to be safe, then
 28 the file may be compromised.
- 29 • Files can often be identified by other attributes in addition to the file name, for example, by comparing file
 30 ownership or creation time. Information regarding a file that has been created and closed can be stored
 31 and then used later to validate the identity of the file when it is reopened. Comparing multiple attributes of
 32 the file improves the likelihood that the file is the expected one.
- 33 • Follow the principle of least privilege when assigning access rights to files.
- 34 • Denying access to a file can prevent an attacker from replacing that file with a link to a sensitive file.
- 35 • Ensure good compartmentalization in the system to provide protected areas that can be trusted.
- 36 • When two or more users, or a group of users, have write permission to a directory, the potential for sharing
 37 and deception is far greater than it is for shared access to a few files. The vulnerabilities that result from
 38 malicious restructuring via hard and symbolic links suggest that it is best to avoid shared directories.
- 39 • Securely creating temporary files in a shared directory is error prone and dependent on the version of the
 40 runtime library used, the operating system, and the file system. Code that works for a locally mounted file
 41 system, for example, may be vulnerable when used with a remotely mounted file system.
- 42 • The mitigation should be centered on converting relative paths into absolute paths and then verifying that
 43 the resulting absolute path makes sense with respect to the configuration and rights or permissions. This
 44 may include checking "whitelists" and "blacklists", authorized super user status, access control lists, or
 45 other fully trusted status.

1 7.15 Missing Required Cryptographic Step [XZS]

2 7.15.1 Description of application vulnerability

3 Cryptographic implementations should follow the algorithms that define them exactly, otherwise encryption can be
4 faulty.

5 7.15.2 Cross reference

6 CWE:

7 325. Missing Required Cryptographic Step

8 7.15.3 Mechanism of failure

9 Not following the algorithms that define cryptographic implementations exactly can lead to weak encryption. This
10 could be the result of many factors such as a programmer missing a required cryptographic step or using weak
11 randomization algorithms.

12 7.15.4 Avoiding the vulnerability or mitigating its effects

13 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 14 • Implement cryptographic algorithms precisely.
- 15 • Use system functions and libraries rather than writing the function.

16 7.16 Insufficiently Protected Credentials [XYM]

17 7.16.1 Description of application vulnerability

18 This weakness occurs when the application transmits or stores authentication credentials and uses an insecure
19 method that is susceptible to unauthorized interception and/or retrieval.

20 7.16.2 Cross reference

21 CWE:

22 256. Plaintext Storage of a Password

23 257. Storing Passwords in a Recoverable Format

24 7.16.3 Mechanism of failure

25 Storing a password in plaintext may result in a system compromise. Password management issues occur when a
26 password is stored in plaintext in an application's properties or configuration file. A programmer can attempt to
27 remedy the password management problem by obscuring the password with an encoding function, such as Base64
28 encoding, but this effort does not adequately protect the password. Storing a plaintext password in a configuration
29 file allows anyone who can read the file access to the password-protected resource. Developers sometimes
30 believe that they cannot defend the application from someone who has access to the configuration, but this attitude
31 makes an attacker's job easier. Good password management guidelines require that a password never be stored
32 in plaintext.

33
34 The storage of passwords in a recoverable format makes them subject to password reuse attacks by malicious
35 users. If a system administrator can recover the password directly or use a brute force search on the information
36 available to him, he can use the password on other accounts.

37 The use of recoverable passwords significantly increases the chance that passwords will be used maliciously. In
38 fact, it should be noted that recoverable encrypted passwords provide no significant benefit over plain-text
39 passwords since they are subject not only to reuse by malicious attackers but also by malicious insiders.

1 7.16.4 Avoiding the vulnerability or mitigating its effects

2 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 3 • Avoid storing passwords in easily accessible locations.
- 4 • Never store a password in plaintext.
- 5 • Ensure that strong, non-reversible encryption is used to protect stored passwords.
- 6 • Consider storing cryptographic hashes of passwords as an alternative to storing in plaintext.

7 7.17 Missing or Inconsistent Access Control [XZN]

8 7.17.1 Description of application vulnerability

9 The software does not perform access control checks in a consistent manner across all potential execution paths.

10 7.17.2 Cross reference

11 CWE:

12 285. Missing or Inconsistent Access Control

13 CERT C guidelines: FIO06-C

14 7.17.3 Mechanism of failure

15 For web applications, attackers can issue a request directly to a page (URL) that they may not be authorized to
16 access. If the access control policy is not consistently enforced on every page restricted to authorized users, then
17 an attacker could gain access to and possibly corrupt these resources.

18 7.17.4 Avoiding the vulnerability or mitigating its effects

19 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 20 • For web applications, make sure that the access control mechanism is enforced correctly at the server side
21 on every page. Users should not be able to access any information simply by requesting direct access to
22 that page, if they do not authorization. Ensure that all pages containing sensitive information are not
23 cached, and that all such pages restrict access to requests that are accompanied by an active and
24 authenticated session token associated with a user who has the required permissions to access that page.

25 7.18 Authentication Logic Error [XZO]

26 7.18.1 Description of application vulnerability

27 The software does not properly ensure that the user has proven their identity.

28 7.18.2 Cross reference

29 CWE:

30 287. Improper Authentication

31 288. Authentication Bypass by Alternate Path/Channel

32 289. Authentication Bypass by Alternate Name

33 290. Authentication Bypass by Spoofing

34 294. Authentication Bypass by Capture-replay

35 301. Reflection Attack in an Authentication Protocol

36 302. Authentication Bypass by Assumed-Immutable Data

37 303. Improper Implementation of Authentication Algorithm

38 305. Authentication Bypass by Primary Weakness

7.18.3 Mechanism of failure

There are many ways that an attacker can potentially bypass the validation of a user. Some of the ways are means of impersonating a legitimate user while others are means of bypassing the authentication mechanisms that are in place. In either case, a user who should not have access to the software system gains access.

Authentication bypass by alternate path or channel occurs when a product requires authentication, but the product has an alternate path or channel that does not require authentication. Note that this is often seen in web applications that assume that access to a particular CGI (Common gateway Interface) program can only be obtained through a "front" screen, but this problem is not just in web applications.

Authentication bypass by alternate name occurs when the software performs authentication based on the name of the resource being accessed, but there are multiple names for the resource, and not all names are checked.

Authentication bypass by capture-replay occurs when it is possible for a malicious user to sniff network traffic and bypass authentication by replaying it to the server in question to the same effect as the original message (or with minor changes). Messages sent with a capture-relay attack allow access to resources that are not otherwise accessible without proper authentication. Capture-replay attacks are common and can be difficult to defeat without cryptography. They are a subset of network injection attacks that rely on listening in on previously sent valid commands, then changing them slightly if necessary and resending the same commands to the server. Since any attacker who can listen to traffic can see sequence numbers, it is necessary to sign messages with some kind of cryptography to ensure that sequence numbers are not simply doctored along with content.

Reflection attacks capitalize on mutual authentication schemes to trick the target into revealing the secret shared between it and another valid user. In a basic mutual-authentication scheme, a secret is known to both a valid user and the server; this allows them to authenticate. In order that they may verify this shared secret without sending it plainly over the wire, they utilize a Diffie-Hellman-style scheme in which they each pick a value, then request the hash of that value as keyed by the shared secret. In a reflection attack, the attacker claims to be a valid user and requests the hash of a random value from the server. When the server returns this value and requests its own value to be hashed, the attacker opens another connection to the server. This time, the hash requested by the attacker is the value that the server requested in the first connection. When the server returns this hashed value, it is used in the first connection, authenticating the attacker successfully as the impersonated valid user.

Authentication bypass by assumed-immutable data occurs when the authentication scheme or implementation uses key data elements that are assumed to be immutable, but can be controlled or modified by the attacker, for example, if a web application relies on a cookie "Authenticated=1".

Authentication logic error occurs when the authentication techniques do not follow the algorithms that define them exactly and so authentication can be jeopardized. For instance, a malformed or improper implementation of an algorithm can weaken the authorization technique.

An authentication bypass by primary weakness occurs when the authentication algorithm is sound, but the implemented mechanism can be bypassed as the result of a separate weakness that is primary to the authentication error.

7.18.4 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Funnel all access through a single choke point to simplify how users can access a resource. For every access, perform a check to determine if the user has permissions to access the resource. Avoid making decisions based on names of resources (for example, files) if those resources can have alternate names.
- Canonicalize the name to match that of the file system's representation of the name. This can sometimes be achieved with an available API (for example, in Win32 the `GetFullPathName` function).
- Utilize some sequence or time stamping functionality along with a checksum that takes this into account to ensure that messages can be parsed only once.
- Use different keys for the initiator and responder or of a different type of challenge for the initiator and responder.

1 7.19 Hard-coded Password [XYP]

2 7.19.1 Description of application vulnerability

3 Hard coded passwords may compromise system security in a way that cannot be easily remedied. It is never a
4 good idea to hardcode a password. Not only does hard coding a password allow all of the project's developers to
5 view the password, it also makes fixing the problem extremely difficult. Once the code is in production, the
6 password cannot be changed without patching the software. If the account protected by the password is
7 compromised, the owners of the system will be forced to choose between security and availability.

8 7.19.2 Cross reference

9 CWE:
10 259. Hard-Coded Password

11 7.19.3 Mechanism of failure

12 The use of a hard-coded password has many negative implications – the most significant of these being a failure of
13 authentication measures under certain circumstances. On many systems, a default administration account exists
14 which is set to a simple default password that is hard-coded into the program or device. This hard-coded password
15 is the same for each device or system of this type and often is not changed or disabled by end users. If a malicious
16 user comes across a device of this kind, it is a simple matter of looking up the default password (which is likely
17 freely available and public on the Internet) and logging in with complete access. In systems that authenticate with a
18 back-end service, hard-coded passwords within closed source or drop-in solution systems require that the back-
19 end service use a password that can be easily discovered. Client-side systems with hard-coded passwords
20 present even more of a threat, since the extraction of a password from a binary is exceedingly simple. If hard-
21 coded passwords are used, it is almost certain that unauthorized users will gain access through the account in
22 question.

23 7.19.4 Avoiding the vulnerability or mitigating its effects

24 Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- 25 • Rather than hard code a default username and password for first time logins, utilize a "first login" mode that
26 requires the user to enter a unique strong password.
- 27 • For front-end to back-end connections, there are three solutions that may be used.
 - 28 1. Use of generated passwords that are changed automatically and must be entered at given time
29 intervals by a system administrator. These passwords will be held in memory and only be valid for
30 the time intervals.
 - 31 2. The passwords used should be limited at the back end to only performing actions for the front end,
32 as opposed to having full access.
 - 33 3. The messages sent should be tagged and checksummed with time sensitive values so as to
34 prevent replay style attacks.

Annex A

(informative)

Guideline Selection Process

A. Guideline Selection Process

It is possible to claim that any language construct can be misunderstood by a developer and lead to a failure to predict program behaviour. A cost/benefit analysis of each proposed guideline is the solution adopted by this Technical Report.

The selection process has been based on evidence that the use of a language construct leads to unintended behaviour (that is, a cost) and that the proposed guideline increases the likelihood that the behaviour is as intended (that is, a benefit). The following is a list of the major source of evidence on the use of a language construct and the faults resulting from that use:

- a list of language constructs having undefined, implementation-defined, or unspecified behaviours,
- measurements of existing source code. This usage information has included the number of occurrences of uses of the construct and the contexts in which it occurs,
- measurement of faults experienced in existing code,
- measurements of developer knowledge and performance behaviour.

The following are some of the issues that were considered when framing guidelines:

- An attempt was made to be generic to particular kinds of language constructs (that is, language independent), rather than being language specific.
- Preference was given to wording that is capable of being checked by automated tools.
- Known algorithms for performing various kinds of source code analysis and the properties of those algorithms (that is, their complexity and running time).

A.1 Cost/Benefit Analysis

The fact that a coding construct is known to be a source of failure to predict correct behaviour is not in itself a reason to recommend against its use. Unless the desired algorithmic functionality can be implemented using an alternative construct whose use has more predictable behaviour, then there is no benefit in recommending against the use of the original construct.

While the cost/benefit of some guidelines may always come down in favor of them being adhered to (for example, don't access a variable before it is given a value), the situation may be less clear-cut for other guidelines. Providing a summary of the background analysis for each guideline will enable development groups.

It is unlikely that all of the guidelines given in this Technical Report will be applicable to all application domains.

A.2 Documenting of the selection process

The intended purpose of this documentation is to enable third parties to evaluate:

- the effectiveness of the process that created each guideline,
- the applicability of individual guidelines to a particular project.

Annex B
(informative)
Template for use in proposing programming language vulnerabilities

B. Skeleton template for use in proposing programming language vulnerabilities

B.1 6.<x> <short title> [<unique immutable identifier>]

No text should appear here—in the space between 6.x and 6.x.1. This is simply an explanation of the header.

Notes on template header. The number "x" depends on the order in which the vulnerabilities are listed in Clause 6. It will be assigned by the editor. The "short title" should be a noun phrase summarizing the description of the application vulnerability. The "unique immutable identifier" is intended to provide an enduring identifier for the vulnerability description, even if their order is changed in the document. No additional text should appear here.

B.1.0 6.<x>.0 Status and history

The header will be removed before publication.

This temporary section will hold the edit history for the vulnerability along with the current status of the vulnerability.

B.1.1 6.<x>.1 Description of application vulnerability

Replace this with a brief description of the application vulnerability. It should be a short paragraph.

B.1.2 6.<x>.2 Cross reference

CWE: Replace this with the CWE identifier. At a later date, other cross-references may be added.

B.1.3 6.<x>.3 Mechanism of failure

Replace this with a brief description of the mechanism of failure. This description provides the link between the programming language vulnerability and the application vulnerability. It should be a short paragraph.

B.1.4 6.<x>.4 Applicable language characteristics

This vulnerability description is applicable to languages with the following characteristics:

Replace this with a bullet list summarizing the pertinent range of characteristics of languages for which this discussion is applicable. This list is intended to assist readers attempting to apply the guidance to languages that have not been treated in the language-specific annexes.

B.1.5 6.<x>.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

Replace this with a bullet list summarizing various ways in which programmers can avoid the programming language vulnerability, break the chain of causation to the application vulnerability, or contain the bad effects of the application vulnerability. Begin with the more direct, concrete, and effective means and then progress to the more indirect, abstract, and probabilistic means.

1 **B.1.6 6.<x>.6 Implications for standardization**

2 *Recommendations for other working groups will be recorded here. For example, we might record suggestions for*
3 *changes to language standards or API standards.*

4 **B.1.7 6.<x>.7 Bibliography**

5 *Insert numbered references for other documents cited in your description. These will eventually be collected into an*
6 *overall bibliography for the TR. So, please make the references complete. Someone will eventually have to*
7 *reformat the references into an ISO-required format, so please err on the side of providing too much information*
8 *rather than too little. Here [1] is an example of a reference:*

9 *[1] Greg Hoglund, Gary McGraw, Exploiting Software: How to Break Code, ISBN-0-201-78695-8, Pearson*
10 *Education, Boston, MA, 2004*

Annex C
(informative)
Template for use in proposing application vulnerabilities

C. Skeleton template for use in proposing application vulnerabilities

C.1 7.<x> <short title> [<unique immutable identifier>]

Notes on template header. The number "x" depends on the order in which the vulnerabilities are listed in Clause 7. It will be assigned by the editor. The "short title" should be a noun phrase summarizing the description of the application vulnerability. The "unique immutable identifier" is intended to provide an enduring identifier for the vulnerability description, even if their order is changed in the document. No additional text should appear here.

C.1.0 7.<x>.0 Status and history

The header will be removed before publication.

This temporary section will hold the edit history for the vulnerability. With the current status of the vulnerability.

C.1.1 7.<x>.1 Description of application vulnerability

Replace this with a brief description of the application vulnerability. It should be a short paragraph.

C.1.2 7.<x>.2 Cross reference

CWE: Replace this with the CWE identifier. At a later date, other cross-references may be added.

C.1.3 7.<x>.3 Mechanism of failure

Replace this with a brief description of the mechanism of failure. It should be a short paragraph.

C.1.4 7.<x>.4 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

Replace this with a bullet list summarizing various ways in which programmers can avoid the application vulnerability, break the chain of causation to the application vulnerability, or contain the bad effects of the application vulnerability. Begin with the more direct, concrete, and effective means and then progress to the more indirect, abstract, and probabilistic means.

C.1.5 7.<x>.5 Bibliography

Insert numbered references for other documents cited in your description. These will eventually be collected into an overall bibliography for the TR. So, please make the references complete. Someone will eventually have to reformat the references into an ISO-required format, so please err on the side of providing too much information rather than too little. Here [1] is an example of a reference:

*[1] Greg Hoglund, Gary McGraw, *Exploiting Software: How to Break Code*, ISBN-0-201-78695-8, Pearson Education, Boston, MA, 2004*

Annex D (informative) Vulnerability Outline

D. Vulnerability Outline

- D.1. Human Factors
 - D.1.1. [BRS] Obscure Language Features
- D.2. Environment
 - D.2.1. [XYN] Adherence to Least Privilege
 - D.2.2. [XYO] Privilege Sandbox Issues
 - D.2.3. Interactions with environment
 - D.2.3.1. [XYS] Executing or Loading Untrusted Code
- D.3. Core Language Issues
 - D.3.1. [BQF] Unspecified Behaviour
 - D.3.2. [EWF] Undefined Behaviour
 - D.3.3. [FAB] Implementation-defined Behaviour
 - D.3.4. [MEM] Deprecated Language Features
 - D.3.5. [BVQ] Unspecified Functionality
- D.4. Pre-processor
 - D.4.1. [NMP] Pre-processor Directives
- D.5. Declarations and Definitions
 - D.5.1. [NAI] Choice of Clear Names
 - D.5.2. [AJN] Choice of Filenames and other External Identifiers
 - D.5.3. [XYR] Unused Variable
 - D.5.4. [YOW] Identifier Name Reuse
 - D.5.4. [BJL] Namespace Issues
- D.6. Types
 - D.6.1. Representation
 - D.6.1.1. [IHN] Type System
 - D.6.1.2. [STR] Bit Representations
 - D.6.1.3. [KLK] Distinguished Values in Data Types
 - D.6.2. Constants
 - D.6.3. Floating-point
 - D.6.3.1. [PLF] Floating-point Arithmetic
 - D.6.4. Enumerated Types
 - D.6.4.1. [CCB] Enumerator Issues
 - D.6.5. Integers
 - D.6.5.1. [FLC] Numeric Conversion Errors
 - D.6.6. Characters and strings
 - D.6.6.1 [CJM] String Termination
 - D.6.7. Arrays
 - D.6.7.1. [XYX] Boundary Beginning Violation
 - D.6.7.2. [XYZ] Unchecked Array Indexing
 - D.6.7.3. [XYW] Unchecked Array Copying
 - D.6.7.4. [XZB] Buffer Overflow
 - D.6.8. Structures and Unions
 - D.6.9. Pointers
 - D.6.9.1. [HFC] Pointer Casting and Pointer Type Changes
 - D.6.9.2. [RVG] Pointer Arithmetic
 - D.6.9.3. [XYH] Null Pointer Dereference
 - D.6.9.4. [XYK] Dangling Reference to Heap
- D.7. Templates/Generics
 - D.7.1. [SYM] Templates and Generics
 - D.7.2. [RIP] Inheritance
- D.8. Initialization
 - D.8.1. [LAV] Initialization of Variables

- 1 D.9. Type Conversions/Limits
- 2 D.9.1. [XYY] Wrap-around Error
- 3 D.9.2. [XZI] Sign Extension Error
- 4 D.10. Operators/Expressions
- 5 D.10.1. [JCW] Operator Precedence/Operator Precedence
- 6 D.10.2. [SAM] Side-effects and Order of Evaluation
- 7 D.10.3. [KOA] Likely Incorrect Expressions
- 8 D.10.4. [XYQ] Dead and Deactivated Code
- 9 D.11. Control Flow
- 10 D.11.1. Conditional Statements
- 11 D.11.1.1. [CLL] Switch Statements and Static Analysis
- 12 D.11.1.2. [EOJ] Demarcation of Control Flow
- 13 D.11.2. Loops
- 14 D.11.2.1. [TEX] Loop Control Variables
- 15 D.11.2.2. [XZH] Off-by-one Error
- 16 D.11.3. Subroutines (Functions, Procedures, Subprograms)
- 17 D.11.3.1. [EWD] Structured Programming
- 18 D.11.3.2. [CSJ] Passing Parameters and Return Values
- 19 D.11.3.3. [DCM] Dangling References to Stack Frames
- 20 D.11.3.4. [OTR] Subprogram Signature Mismatch
- 21 D.11.3.5. [GDL] Recursion
- 22 D.11.3.7. [NZN] Returning Error Status
- 23 D.11.4. Termination Strategy
- 24 D.11.4.1. [REU] Termination Strategy
- 25 D.12. Internal interfaces
- 26 D.12.1 [LRM] Extra Intrinsic
- 27 D.13. External interfaces
- 28 D.13.1. Memory Management
- 29 D.13.1.1. [AMV] Type-breaking Reinterpretation of Data
- 30 D.13.1.2. [XYL] Memory Leak
- 31 D.13.1.3. [XZX] Memory Locking
- 32 D.13.1.4. [XZP] Resource Exhaustion
- 33 D.13.2. Input
- 34 D.13.2.1. [RST] Injection
- 35 D.13.2.2. [XYT] Cross-site Scripting
- 36 D.13.2.3. [XZQ] Unquoted Search Path or Element
- 37 D.13.2.4. [XZR] Improperly Verified Signature
- 38 D.13.2.5. [XZL] Discrepancy Information Leak
- 39 D.13.3. Output
- 40 D.13.3.1. [XZK] Sensitive Information Uncleared Before Use
- 41 D.13.4. Libraries
- 42 D.13.4.1. [TRJ] Argument Passing to Library Functions
- 43 D.13.4.2. [NYY] Dynamically-linked Code and Self-modifying Code
- 44 D.13.4.3. [NSQ] Library Signature
- 45 D.13.4.4. [HJW] Unanticipated Exceptions from Library Routines
- 46 D.13.5. Files
- 47 D.13.5.1. [EWR] Path Traversal
- 48 D.14. Miscellaneous
- 49 D.14.1. [XZS] Missing Required Cryptographic Step
- 50 D.14.2. Authentication
- 51 D.14.2.1. [XYM] Insufficiently Protected Credentials
- 52 D.14.2.2. [XZN] Missing or Inconsistent Access Control
- 53 D.14.2.3. [XZO] Authentication Logic Error
- 54 D.14.2.4. [XYP] Hard-coded Password

Annex E (informative) Language Specific Vulnerability Template

E. Vulnerability descriptions for <language>

E.1 <language>.1 Identification of standards

[This section should list the relevant language standards and other documents that describe the language treated in the annex. It should not be simply a list of standards. It should do whatever is required to describe the language that is the baseline. In some cases, it might be a standard plus some other documents, or a standard minus the annex that lists deprecated features. It might include some explanation, such as "don't use any features that are undefined".]

E.2 <language>.2 General terminology and concepts

[This section should provide an overview of general terminology and concepts that are utilized throughout the annex.]

Every vulnerability description of Clause 6 of the main document should be addressed in the annex in the same order even if there is simply a notation that it is not relevant to the language in question.

Each vulnerability description should have the following format:

E.3 <language>.<x> <Vulnerability Name> [<3 letter tag>]

<language>.<x>.0 Status and history

[Revision history. This section will eventually be removed.]

<language>.<x>.1 Terminology and features

[In this and other sections, if there is nothing to be explained, simply say "None".]

[This section should describe terms that are in the language standard and which are used in the explanation that follows.]

term: An explanation in the form of one or more complete sentences.

<language>.<x>.2 Description of vulnerability

[This merges the prior sections for description and mechanism. Examples, both good and bad, are strongly encouraged.]

<language>.<x>.3 Avoiding the vulnerability or mitigating its effects

- [An imperative sentence followed by optional additional sentences written in the indicative.]
- ...

<language>.<x>.4 Implications for standardization

Future standardization efforts should consider:

- Requiring ...
- Adding ...

- 1 • Changing ...
2 • Other verbs ending in "ing"
3

4 In those cases where a vulnerability is simply not applicable to the language, the following format should be
5 used:

6 **<language>.<x> <Vulnerability Name> [<3 letter tag>]**

7 This vulnerability is not applicable to <language>. [Optionally, an explanation of inapplicability may be added,
8 including qualifications and pointers to other related vulnerabilities that might be present.]

9

Bibliography

- [1] ISO/IEC Directives, Part 2, *Rules for the structure and drafting of International Standards*, 2001
- [2] ISO/IEC TR 10000-1, *Information technology — Framework and taxonomy of International Standardized Profiles — Part 1: General principles and documentation framework*
- [3] ISO 10241, *International terminology standards — Preparation and layout*
- [4] ISO/IEC TR 15942:2000, "Information technology - Programming languages - Guide for the use of the Ada programming language in high integrity systems"
- [5] Joint Strike Fighter Air Vehicle: C++ Coding Standards for the System Development and Demonstration Program. Lockheed Martin Corporation. December 2005.
- [6] ISO/IEC 15291:1999, Information technology - Programming languages - Ada Semantic Interface Specification (ASIS)
- [7] Software Considerations in Airborne Systems and Equipment Certification. Issued in the USA by the Requirements and Technical Concepts for Aviation (document RTCA SC167/DO-178B) and in Europe by the European Organization for Civil Aviation Electronics (EUROCAE document ED-12B). December 1992.
- [8] IEC 61508: Parts 1-7, Functional safety: safety-related systems. 1998. (Part 3 is concerned with software).
- [9] ISO/IEC 15408: 1999 Information technology. Security techniques. Evaluation criteria for IT security.
- [10] J Barnes, High Integrity Software - the SPARK Approach to Safety and Security. Addison-Wesley. 2002.
- [11] R. Seacord, The CERT C Secure Coding Standard. Boston, MA: Addison-Wesley, 2008.
- [12] Motor Industry Software Reliability Association. *Guidelines for the Use of the C Language in Vehicle Based Software*, 2004 (second edition)².
- [13] ISO/IEC TR24731-1, *Extensions to the C Library, — Part I: Bounds-checking interfaces*
- [14] Steve Christy, *Vulnerability Type Distributions in CVE*, V1.0, 2006/10/04
- [15] *ARIANE 5: Flight 501 Failure*, Report by the Inquiry Board, July 19, 1996
<http://esamultimedia.esa.int/docs/esa-x-1819eng.pdf>
- [16] Motor Industry Software Reliability Association. *Guidelines for the Use of the C++ Language in critical systems*, June 2008
- [17] ISO/IEC TR 24718: 2004, *Guide for the use of the Ada Ravenscar Profile in high integrity systems*
- [18] L. Hatton, Safer C: developing software for high-integrity and safety-critical systems. McGraw-Hill 1995
- [19] Robert W. Sebesta, Concepts of Programming Languages, 8th edition, ISBN-13: 978-0-321-49362-0, ISBN-10: 0-321-49362-1, Pearson Education, Boston, MA, 2008

² The first edition should not be used or quoted in this work.

- 1 [20] Hogaboom, Richard, *A Generic API Bit Manipulation in C*, Embedded Systems Programming, Vol 12, No 7,
2 July 1999 <http://www.embedded.com/1999/9907/9907feat2.htm>
- 3 [21] Carlo Ghezzi and Mehdi Jazayeri, *Programming Language Concepts*, 3rd edition, ISBN-0-471-10426-4,
4 John Wiley & Sons, 1998
- 5 [22] Lions, J. L. [ARIANE 5 Flight 501 Failure Report](#)²⁵. Paris, France: European Space Agency (ESA) & National
6 Center for Space Study (CNES) Inquiry Board, July 1996.
- 7 [23] Seacord, R. *Secure Coding in C and C++*. Boston, MA: Addison-Wesley, 2005. See
8 <http://www.cert.org/books/secure-coding>²⁶ for news and errata.
- 9 [24] John David N. Dionisio. Type Checking. <http://myweb.lmu.edu/dondi/share/pl/type-checking-v02.pdf>
- 10 [25] MISRA Limited. "[MISRA C](#)²⁷: 2004 Guidelines for the Use of the C Language in Critical Systems."
11 Warwickshire, UK: MIRA Limited, October 2004 (ISBN 095241564X).
- 12 [26] The Common Weakness Enumeration (CWE) Initiative, MITRE Corporation, (<http://cwe.mitre.org/>)
- 13 [27] Goldberg, David, *What Every Computer Scientist Should Know About Floating-Point Arithmetic*, ACM
14 Computing Surveys, vol 23, issue 1 (March 1991), ISSN 0360-0300, pp 5-48.
- 15 [28] IEEE Standards Committee 754. IEEE Standard for Binary Floating-Point Arithmetic, ANSI/IEEE Standard
16 754-2008. Institute of Electrical and Electronics Engineers, New York, 2008.
- 17 [29] Bo Einarsson, ed. *Accuracy and Reliability in Scientific Computing*, SIAM, July 2005
18 <http://www.nsc.liu.se/wg25/book>
- 19 [30] GAO Report, *Patriot Missile Defense: Software Problem Led to System Failure at Dhahran, Saudi Arabia*,
20 B-247094, Feb. 4, 1992, <http://archive.gao.gov/t2pbat6/145960.pdf>
- 21 [31] Robert Skeel, *Roundoff Error Cripples Patriot Missile*, SIAM News, Volume 25, Number 4, July 1992, page
22 11, <http://www.siam.org/siamnews/general/patriot.htm>
- 23 [32] CERT. *CERT C++ Secure Coding Standard*.
24 <https://www.securecoding.cert.org/confluence/pages/viewpage.action?pageId=637> (2009).
- 25 [33] P. V. Bhansali, A systematic approach to identifying a safe subset for safety-critical software, ACM
26 SIGSOFT Software Engineering Notes, v.28 n.4, July 2003
- 27 [34] Ada 95 Quality and Style Guide, SPC-91061-CMC, version 02.01.01. Herndon, Virginia: Software
28 Productivity Consortium, 1992. Available from: <http://www.adaic.org/docs/95style/95style.pdf>
- 29 [35] Ghassan, A., & Alkadi, I. (2003). Application of a Revised DIT Metric to Redesign an OO Design. *Journal of*
30 *Object Technology* , 127-134.
- 31 [36] Subramanian, S., Tsai, W.-T., & Rayadurgam, S. (1998). Design Constraint Violation Detection in Safety-
32 Critical Systems. The 3rd IEEE International Symposium on High-Assurance Systems Engineering , 109 -
33 116.