



MISRA C:2023

Guidelines for the use of the C language in critical systems

Third edition, Second revision

April 2023



First published April 2023 by The MISRA Consortium Limited
1 St James Court
Whitefriars
Norwich
Norfolk
NR3 1RU
UK

www.misra.org.uk

Copyright © 2023 The MISRA Consortium Limited

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical or photocopying, recording or otherwise without the prior written permission of the Publisher.

“MISRA”, “MISRA C” and the triangle logo are registered trademarks owned by The MISRA Consortium Limited.

Other product or brand names are trademarks or registered trademarks of their respective holders and no endorsement or recommendation of these products by MISRA is implied.

ISBN 978-1-911700-08-1 paperback
ISBN 978-1-911700-09-8 PDF

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

License terms: This copy of MISRA C:2023 Third Edition, Second Revision is licensed to CodeSecure, Inc. (hereafter referred to as “The Licensee”) according to the terms of the License Agreement dated 10 Oct 2023 and the purposes stated therein. Specifically this document is licensed for distribution to licensed users of the Licensee’s product CodeSonar.

The Licensee may not change, modify, amend or develop this document in any way without the prior written consent of MISRA.

No permission is given for use or distribution of this document by or to individuals or companies who are not employees of The Licensee or who are not licensed users of CodeSonar.

For full details of the license terms please refer to your License Agreement. You agree to be bound by these license terms when using this document.

MISRA C:2023

Guidelines for the use of the C language in critical systems

Third edition, Second revision

April 2023

MISRA Mission Statement

We provide world-leading, best practice guidelines for the safe and secure application of both embedded control systems and standalone software.

MISRA is a collaboration between manufacturers, component suppliers and engineering consultancies which seeks to promote best practice in developing safety- and security-related electronic systems and other software-intensive applications. To this end, MISRA publishes documents that provide accessible information for engineers and management, and holds events to permit the exchange of experiences between practitioners.

Disclaimer

Adherence to the requirements of this document does not in itself ensure error-free robust software or guarantee portability and re-use.

Compliance with the requirements of this document, or any other standard, does not of itself confer immunity from legal obligations.

Foreword

In April 1998, the sterling efforts of a small group of software/systems engineers resulted in the publication of the first edition of MISRA-C [1].

Few, if any, would have anticipated the impact that that document was to have had, nor the reach that it would attain — and 25 years on (through a second edition in 2004 [2] and the third edition in 2012 [3]) is still having and is still attaining.

To mark that silver anniversary, this second revision to that third edition of the MISRA C Guidelines consolidates six updates to MISRA C:2012 [3], reflecting ongoing work within the Working Group:

- MISRA C:2012 Amendment 1, *Additional security guidelines for MISRA C:2012* [5]
- MISRA C:2012 Amendment 2, *Updates for ISO/IEC 9899:2011 Core functionality* [6]
- MISRA C:2012 Amendment 3, *Updates for ISO/IEC 9899:2011/2018 Phase 2 — New C11/C18 features* [7]
- MISRA C:2012 Amendment 4, *Updates for ISO/IEC 9899:2011/2018 Phase 3 — Multi-threading and atomics* [8]
- MISRA C:2012 Technical Corrigendum 1 [9]
- MISRA C:2012 Technical Corrigendum 2 [10]

Note: The first revision [4] consolidated MISRA C:2012 [3] with AMD1 [5] and TC1 [9]

These enhancements to MISRA C:2012 introduce support for the new features introduced by C11 and C18, as well as additional guidance on existing language features.

This revision also makes the enhanced guidance on achieving compliance with MISRA C, as provided by MISRA Compliance:2020 [11] a mandatory part of the development process.

We trust that this revision to MISRA C will be welcomed by the community at large, and will offer confidence to projects and organizations who have held off migrating to C11 or C18, as well as continuing to provide best practice for earlier editions of the C Standard.

Andrew Banks MIET FBCS CITP
Chairman, MISRA C Working Group

Acknowledgements

The MISRA C Working Group

The MISRA Consortium would like to thank the following current members of the MISRA C Working Group, and their employers, for their significant contribution to the writing of this document:

Roberto Bagnara	BUGSENG (and the University of Parma)
Dave Banham	Blackberry Ltd
Andrew Banks	LDRA Ltd (also Intuitive Consulting)
Jill Britton	Perforce Software Inc.
Alex Gilding	Perforce Software Inc.
Daniel Kästner	AbsInt Angewandte Informatik GmbH
Gerlinde Kettl	Vitesco Technologies GmbH
Michal Rozenau	Parasoft Corp.
Chris Tapp	LDRA Ltd (also Keylevel Consultants Ltd)

The MISRA Consortium also wishes to acknowledge contributions from the following individual during the review process for this document:

David Ward	HORIBA MIRA Ltd
------------	-----------------

The MISRA Consortium would like to thank the following past members of the MISRA C Working Group for their various contributions to this document:

Fulvio Baccaglini	Mark Bradbury	Paul Burden	Mark Dawson-Butterworth
Ben Godwood	Mike Hennell	Chris Hills	Gavin McCall
Steve Montgomery	Thomas M. Tuerke	Liz Whiting	

Other Acknowledgements

DokuWiki was used extensively during the drafting of this document. Our thanks go to all those involved in its development.

This document was typeset using Open Sans. Open Sans is a trademark of Google and may be registered in certain jurisdictions. Digitized data copyright © 2010–2011, Google Corporation. Licensed under the Apache License, Version 2.0.

The descriptions of implementation-defined behaviours in Appendix G have been reproduced from editions of the C Standard (BS ISO/IEC 9899) published by BSI Standards Limited; the text is identical to that in the equivalent ISO editions. Permission to reproduce extracts from British Standards is granted by the BSI Standards Limited (BSI) under Licence No. 2013ET0003. No other use of this material is permitted.

British Standards can be obtained in PDF or hard copy formats from the BSI online shop: www.bsigroup.com/Shop or by contacting BSI Customer Services for hard copies only: Tel: +44 20 8996 9001, Email: cservices@bsigroup.com.

Contents

1	Introduction	1
	1.1 Background	1
	1.2 The vision	1
	1.3 Scope	1
	1.4 Adoption	2
2	Background to MISRA C	3
	2.1 The popularity of C	3
	2.2 Disadvantages of C	3
	2.3 The use of C in critical systems	4
3	Tool selection	5
4	Prerequisite knowledge	6
5	Adopting and using MISRA C	7
6	Introduction to the guidelines	8
	6.1 Guideline classification	8
	6.2 Guideline categories	8
	6.3 Organization of guidelines	9
	6.4 Redundancy in the guidelines	9
	6.5 Decidability of rules	9
	6.6 Scope of analysis	10
	6.7 Multi-organization projects	11
	6.8 Automatically generated code	12
	6.9 Presentation of guidelines	13
	6.10 Understanding the source references	15
7	Directives	17
	7.1 The implementation	17
	7.2 Compilation and build	19
	7.3 Requirements traceability	19
	7.4 Code design	20
	7.5 Concurrency considerations	37
8	Rules	42
	8.1 A standard C environment	42
	8.2 Unused code	46
	8.3 Comments	52
	8.4 Character sets and lexical conventions	54
	8.5 Identifiers	55

8.6	Types	66
8.7	Literals and constants	68
8.8	Declarations and definitions	75
8.9	Initialization	91
8.10	The essential type model	99
8.11	Pointer type conversions	112
8.12	Expressions	124
8.13	Side effects	131
8.14	Control statement expressions	139
8.15	Control flow	146
8.16	Switch statements	154
8.17	Functions	160
8.18	Pointers and arrays	170
8.19	Overlapping storage	183
8.20	Preprocessing directives	185
8.21	Standard libraries	196
8.22	Resources	219
8.23	Generic Selections	239
9	References	250
9.1	MISRA Publications	250
9.2	The C Standard	251
9.3	Other International Standards	252
9.4	Other References	252
Appendix A	Summary of guidelines	254
Appendix B	Guideline attributes	266
Appendix C	Type safety issues with C	272
Appendix D	Essential types	276
Appendix E	Applicability to automatically generated code	283
Appendix F	Obsolescent language features	287
Appendix G	Implementation-defined behaviour checklist	288
Appendix H	Undefined and critical unspecified behaviour	292
Appendix I	Example deviation record	303
Appendix J	Glossary	304
Appendix K	Change log	309

1 Introduction

1.1 Background

The MISRA C Guidelines define a subset of the C language in which the opportunity to make mistakes is either removed or reduced. Many standards for the development of safety-related software require, or recommend, the use of a language subset, and this can also be used to develop any application with security, high integrity or high reliability requirements.

As well as defining this subset, these MISRA C Guidelines will:

- Provide educational material for those developing C programs;
- Provide reference material for tool developers.

1.2 The vision

It is a long-term objective for MISRA C to define a “predictable subset” of the C language, and to provide explicit guidance for the avoidance of all instances of undefined and unspecified behaviour. For this reason, features that are initially permitted without the support of specific guidelines may be subject to restrictions in the future.

The vision for MISRA C is to:

- Enhance the existing guidance:
 - Correct any known issues with the previous editions (where still relevant);
 - Add new guidelines for which there is a strong rationale;
 - Improve the specification and the rationale for existing guidelines, where appropriate;
 - Remove any guidelines for which the rationale is insufficient;
 - Increase the number of guidelines that can be processed by static analysis tools, by improving the decidability where possible;
- Provide guidance on the applicability of the guidelines to automatically-generated code.

1.3 Scope

The editions of the C Standard supported by this document are:

Edition		Amendments and Corrections	
ISO/IEC 9899:1990	[19]	referred to as “C90”	as revised by [20], [21], [22]
ISO/IEC 9899:1999	[23]	referred to as “C99”	as revised by [24], [25], [26]
ISO/IEC 9899:2011	[27]	referred to as “C11”	as revised by [28]
ISO/IEC 9899:2018	[29]	referred to as “C18”	

Notes:

1. Access to a copy of the relevant C Standard is not necessary for the use of MISRA C, but it may be helpful.
2. ISO/IEC 9899:2018 [29] does not introduce any functional changes. As such, any references to C11 are equally applicable to C18.
3. The edition of the C Standard chosen may be influenced by factors such as the amount of legacy code being reused within a project and/or the availability of compilers for the target environment.

1.4 Adoption

MISRA C should be adopted at the outset of a project.

In addition, the guidance given in MISRA Compliance [11] shall be followed when making a claim of compliance.

Note: If a project is building on existing code that has a proven track record, then the benefits of compliance with MISRA C may be outweighed by the risks of introducing defects when making the code compliant. In such cases, a judgement should be made between the benefits of adoption and the risks of introducing defects.

2 Background to MISRA C

2.1 The popularity of C

The C programming language is popular because:

- C compilers are readily available for many processors;
- C programs can be compiled to efficient machine code;
- It is defined by an international standard;
- It provides mechanisms to access the input/output capabilities of the target processor, whether directly or by means of language extensions;
- There is a considerable body of experience with using C in critical systems;
- It is widely supported by static analysis and test tools.

2.2 Disadvantages of C

While popular, the language has several drawbacks which are discussed in the following sub-sections.

2.2.1 Language definition

The C Standard does not specify the language completely but places some aspects under the control of an implementation. This is intentional, partly because of the desire to support many pre-existing implementations for widely different target processors.

As a result there are areas of the language in which:

- The behaviour is undefined;
- The behaviour is unspecified;
- An implementation is free to choose its own behaviour provided that it is documented.

A program that relies on undefined or unspecified behaviour is not necessarily guaranteed to behave in a predictable manner.

A program that places excessive reliance on implementation-defined behaviour may be difficult to port to a different target. The presence of implementation-defined behaviour may also hinder static analysis if it is not possible to configure the analyser to handle it.

2.2.2 Language misuse

While C programs can be laid out in a structured and comprehensible manner, C makes it easy for programmers to write obscure code that is difficult to understand.

The specification of the operators makes it difficult for programming errors to be detected by a compiler. For example, the following two fragments of code are both perfectly legal so it is impossible for a compiler to know whether one has been mistakenly used in place of the other:

```
if ( a == b )    /* tests whether a and b are equal          */
if ( a = b )    /* assigns b to a and tests whether a is non-zero */
```

2.2.3 Language misunderstanding

There are areas of the language that are commonly misunderstood by programmers. For example, C has more operators than some other languages and consequently has a high number of different operator precedence levels, some of which are not intuitive.

The type rules provided by C can also be confusing to programmers who are familiar with strongly-typed languages. For example, operands may be “promoted” to wider types, meaning that the type resulting from an operation is not necessarily the same as that of the operands.

2.2.4 Run-time error checking

C programs can be compiled into small and efficient machine code, but the trade-off is that there is a very limited degree of run-time checking. C programs generally do not provide run-time checking for common problems such as arithmetic exceptions (e.g. divide by zero), overflow, validity of pointers, or array bound errors. The C philosophy is that the programmer is responsible for making such checks explicitly.

2.3 The use of C in critical systems

Notwithstanding the continued use of Ada and the emergence of Rust, the C programming language is the *de facto* language choice for software in critical systems, whether safety- or security-related, for both embedded (freestanding) and hosted applications.

The recommendations within these Guidelines, when used within a documented software development process, address many of the disadvantages of the C language, irrespective of the purpose of the end-use of the developed code.

These Guidelines are therefore equally as applicable within a security-related environment as they are within a safety-related one.

3 Tool selection

Withdrawn — superseded by Sections 2.6.1 and 2.6.2 of MISRA Compliance:2020 [11].

4 Prerequisite knowledge

Withdrawn — superseded by Sections 2.3, 2.6.4, 2.6.5 and 7.1 of MISRA Compliance:2020 [11].

5 Adopting and using MISRA C

Withdrawn — superseded by MISRA Compliance:2020 [11].

6 Introduction to the guidelines

This section explains the presentation of the guidelines in Section 7 and Section 8, the main content of this document, and serves as an introduction to those sections.

6.1 Guideline classification

Every MISRA C guideline is classified as either being a “rule” or a “directive”.

A directive is a guideline for which it is not possible to provide the full description necessary to perform a check for compliance. Additional information, such as might be provided in design documents or requirements specifications, is required in order to be able to perform the check. Static analysis tools may be able to assist in checking compliance with directives but different tools may place widely different interpretations on what constitutes a non-compliance.

A rule is a guideline for which a complete description of the requirement has been provided. It should be possible to check that source code complies with a rule without needing any other information. In particular, static analysis tools should be capable of checking compliance with rules subject to the limitations described in Section 6.5.

Section 7 contains all the directives and Section 8 contains all the rules.

6.2 Guideline categories

Every MISRA C guideline is given a single category of “mandatory”, “required” or “advisory”, whose meanings are described below.

Beyond this basic classification the document does not give, nor intend to imply, any grading of importance of each of the guidelines, whether rules or directives. All mandatory guidelines should be considered to be of equal importance, as should all required guidelines and all advisory guidelines.

6.2.1 Mandatory guidelines

C code which is claimed to conform to this document shall comply with every mandatory guideline — deviation from mandatory guidelines is not permitted.

Note: if a checking tool produces a diagnostic message, this does not necessarily mean that a guideline has been violated for the reasons given in Section 6.5.

6.2.2 Required guidelines

C code which is claimed to conform to this document shall comply with every required guideline, with a formal deviation required, as described in Section 4 of MISRA Compliance:2020 [11], where this is not the case.

An organization or project may choose to treat any required guideline as if it were mandatory.

6.2.3 Advisory guidelines

These are recommendations. However, the status of “advisory” does not mean that these items can be ignored, but rather that they should be followed as far as is reasonably practical. Formal deviation is not necessary for advisory guidelines but, if the formal deviation process is not followed, alternative arrangements should be made for documenting non-compliances.

An organization or project may choose to treat any advisory guideline as if it were mandatory or required.

6.3 Organization of guidelines

The guidelines are organized under different topics within the C language. However there is inevitably overlap, with one guideline possibly being relevant to a number of topics. Where this is the case the guideline has been placed under the most relevant topic.

6.4 Redundancy in the guidelines

There are a few cases within this document where a guideline is given that refers to a language feature that is banned or advised against elsewhere in the document. This is intentional. It may be that the user chooses to use that feature, either by raising a deviation against a required guideline, or by choosing not to follow an advisory guideline. In this case the second guideline, constraining the use of that feature, becomes relevant.

6.5 Decidability of rules

Each mandatory, required and advisory rule is classified as *decidable* or *undecidable*. This classification describes the theoretical ability of a static analyser to answer the question “Does this code comply with this rule?” The directives are **not** classified in this way because it is impossible, given only the source code, to devise an algorithm that could guarantee to check for compliance.

A rule is *decidable* if it is possible for a program to answer the question with a “yes” or a “no” **in every case** and *undecidable* otherwise. A review of the theory of computation, on which this classification is based, is beyond the scope of this document but a rule is likely to be undecidable if detecting violations depends on run-time properties such as:

- The value that an object holds;
- Whether control reaches a particular point in the program.

Decidable rules have useful properties with regard to static analysis. Provided that a defect-free and complete static analyser is configured correctly:

- A reported violation of a decidable rule indicates a real violation;
- No reported violation of a decidable rule indicates there are no violations in the code being analysed.

Some examples of decidable rules are:

- Rule 5.2: depends on the names and scopes of identifiers;
- Rule 11.3: depends on the source pointer and destination pointer types;
- Rule 20.7: depends on the syntactic form of the result of a macro expansion.

Static analysers vary in their ability to detect violations of undecidable rules:

- A reported violation of an undecidable rule may not necessarily indicate a real violation; some analysers take the approach of reporting **possible** violations to remind users of the uncertainty;
- No reported violation of an undecidable rule does not necessarily indicate that there are no violations in the code being analysed.

Some examples of undecidable rules are:

- Rule 12.2: depends on the value of the right-hand operand of a shift operator;
- Rule 2.1: depends on knowing whether execution never reaches a certain point.

As indicated in Section 3.4 of MISRA Compliance:2020 [11], a process should be developed for analysing the results of static analysis and recording the outcome. Particular attention should be paid to the process for analysing any output that relates to undecidable rules.

6.6 Scope of analysis

Each rule is classified according to the amount of code that needs to be checked in order to detect violations. As for decidability, the concept of analysis scope is not applied to directives.

The analysis scopes that may be applied to rules are “Single Translation Unit” and “System”.

If a rule is classified as capable of being checked on a “Single Translation Unit” basis then it is possible to detect all violations within a project by checking each translation unit independently. For example, the presence of *switch* statements that do not contain *default* labels (Rule 16.4) within one translation unit has no effect on whether other translation units contain such *switch* statements.

If a rule is classified as needing to be checked on a “System” basis then identifying violations of a rule within a translation unit requires checking more than the translation unit in question. Rules that are classified as “System” are best checked by analysing all the source code, although it will be possible to identify some violations when checking a subset of the whole source. For example, if a project has two translation units *A* and *B*, it is possible to check that all declarations and definitions of an object within each translation unit use the same type names and qualifiers (Rule 8.3). However, this does not guarantee that the declarations and definitions in *A* use the same type names and qualifiers as those in *B*. All of the source code that will be compiled and linked into the executable therefore needs to be checked to guarantee compliance with this rule.

All undecidable rules need to be checked on a “System” basis because, in the general case, information about the behaviour of other translation units will be needed. For example, whether or not the value of the automatic object *x* is set before *x* is used (Rule 9.1) in the function *g*, below, will depend on the behaviour of the function *f* which is defined in another translation unit:

```
extern void f ( uint16_t *p );

uint16_t y;

void g ( void )
{
    uint16_t x; /* x is not given a value */

    f ( &x ); /* f might modify the object pointed to by its parameter */
    y = x; /* x may or may not be unset */
}
```

6.7 Multi-organization projects

Projects may involve code from various organizations, for example:

- The Standard Library from the compiler implementer;
- Low-level driver code from a device vendor;
- Operating system and high-level driver code from a specialist supplier;
- Application code which may be shared between collaborating organizations according to their expertise.

The Standard Library is likely to be concerned with efficiency of execution. It may rely on implementation details or unspecified behaviours such as:

- Casting pointers to object types into pointers to other object types;
- The layout of stack frames, and in particular the locations of function parameters within those frames;
- Embedding assembly language statements in C.

As it is part of the implementation, and its functionality and interface is defined in the C Standard, the Standard Library is not required to comply with MISRA C. Unless otherwise specified in the individual guidelines, the contents of standard *header files*, and any files that are included during processing of a standard *header file*, are not required to comply with MISRA C. However, guidelines that rely on the interface provided by standard header declarations and macros are still applicable. For example, the *essential type* rules apply to the types of arguments passed to functions specified in the Standard Library and to their results.

All other code should comply with the MISRA C guidelines to the greatest extent possible. If all the source code is available then the entire program can be checked for compliance. However, in many projects it is not possible for the developing organization to obtain all the source code. For example, collaborating organizations may share functional specifications, interface specifications and object code but are likely to protect their own intellectual property by not sharing source code. When only part of the source code is available for analysis, other techniques should be applied in order to confirm compliance. For example:

- Organizations that supply commercial code may be willing to issue a statement of MISRA C compliance;
- Organizations that are collaborating might:
 - Agree upon a common procedure and tools that each will apply to their own source code;
 - Supply each other with stub versions of source code to permit cross-organizational checks to be made.

A project should define which of the following methods it will use to deal with code supplied by other organizations:

1. In the same manner as code being developed in-house — this relies on having the source code available;
2. In the same manner as the Standard Library — this relies on the organization supplying a suitable MISRA C compliance statement;
3. Compliance checks to be performed on the contents of *header files* and their interface — this assumes that no compliance statement is available and no source code, other than *header files*, is available.

In case (3), appropriate additional verification and validation techniques should be employed prior to using the code.

Note: some process standards may include wider requirements for managing multi-organization developments, for example ISO 26262 [34] Part 8 Clause 5 “Interfaces in distributed development”.

6.8 Automatically generated code

The MISRA C guidelines are applicable to code that has been generated automatically. Responsibility for compliance lies both with the developer of the automatic code generation tool, and with the developer of the model from which code is being generated. Since there are several modelling packages, each of which may have several automatic code generators, it is not possible to allocate this responsibility individually for each MISRA C guideline. It is expected that users of modelling packages and code generators will employ relevant guidelines such as MISRA AC GMG [13], MISRA AC SLSF [14] and MISRA AC TL [15].

The category assigned to a MISRA C guideline when applied to automatically generated code is not necessarily the same as that when applied to manually generated code. When making this distinction, it is important to note that some modelling packages permit the injection of manually generated code into a model. Such code is **not** treated as having been generated automatically when determining the applicable category of a guideline.

Appendix E identifies those MISRA C guidelines whose category as applied to automatically generated code is different from that for manually generated code. It also states the documentation requirements that are placed on automatic code generator developers.

6.9 Presentation of guidelines

The individual requirements of this document are presented in the following format:

Ident	Requirement text	[Source ref]
Category	Category	
Analysis	Decidability, Scope	
Applies to	Cxx	

where:

- "Ident" is a unique identifier for the guideline;
- "Requirement text" is the guideline itself;
- "Source ref" indicates the primary source(s) which led to this item or group of items, where applicable. See Section 6.10 for an explanation of the significance of these references, and a key to the source materials;
- "Category" is one of "Mandatory", "Required" or "Advisory", as explained in Section 6.2;
- "Decidability" is one of "Decidable" or "Undecidable", as explained in Section 6.5;
- "Scope" is one of "System" or "Single Translation Unit", as explained in Section 6.6;
- "Cxx" is a comma separated list indicating which edition(s) of the C Standard that the guideline applies to.

Notes:

1. Where a guideline does not apply to the chosen edition of the C Standard, it is treated as "not applicable" for the purposes of MISRA Compliance [11].
2. Any guideline applying to C11 also applies to C18.
3. Since directives do not have a decidability or a scope of analysis, the "Analysis" line is omitted for directives.

In addition, supporting text is provided for each item or group of related items. The text gives, where appropriate, some explanation of the underlying issues being addressed by the guideline(s), and examples of how to apply the guideline(s).

Within the supporting text, there may be a heading titled "Amplification", followed by text that provides a more precise description of the guideline. An amplification is normative; if it conflicts with the headline, the amplification takes precedence. This mechanism is convenient as it allows a complicated concept to be conveyed using a short headline.

Within the supporting text, there may be a heading titled "Exception", followed by text that describes situations in which the rule does not apply. The use of exceptions permits the description of some guidelines to be simplified. It is important to note that an exception is a situation in which the guideline does not apply. Code that complies with a guideline by virtue of an exception does not require a deviation.

Within the supporting text, there may be a heading titled “Example”, followed by code snippets demonstrating the application of the guideline. These code snippets may be incomplete, for the sake of brevity (for example, an *if* statement without its body, or the omission of function call return value checking).

Within the supporting text, there may be a heading titled “See also”, followed by a list of other guidelines which are related to or interact with the guideline.

The supporting text is not intended as a tutorial in the relevant language feature, as the reader is assumed to have a working knowledge of the language. Further information on the language features can be obtained by consulting the relevant section of the C Standard or other C language reference books. Where a source reference is given for one or more of the *portability issues* listed in the C Standard, then the original issue raised in it may provide additional help in understanding the guideline.

Within the guidelines, and their supporting text, the following font styles are used to represent C keywords and C code:

- C keywords appear in *italic text*;
- Items defined in the Glossary appear in *italic text*;
- C code appears in a monospaced font, either within other text;
- Or

as separate code fragments.

In code fragments the following *typedef*d types have been assumed, to comply with Dir 4.6:

```
uint8_t           /* unsigned 8-bit integer      */
uint16_t          /* unsigned 16-bit integer     */
uint32_t          /* unsigned 32-bit integer     */
int8_t            /* signed 8-bit integer        */
int16_t           /* signed 16-bit integer       */
int32_t           /* signed 32-bit integer       */

float32_t         /* 32-bit floating-point - real */
float64_t         /* 64-bit floating-point - real */
float128_t        /* 128-bit floating-point - real */
cfloat32_t        /* 32-bit floating-point - complex */
cfloat64_t        /* 64-bit floating-point - complex */
cfloat128_t       /* 128-bit floating-point - complex */
```

Non-specific object names are constructed to give an indication of the type. For example:

```
uint8_t    u8a;    /* 8-bit unsigned integer      */
int16_t    s16b;   /* 16-bit signed integer       */
int32_t    s32c;   /* 32-bit signed integer       */

bool_t     bla;    /* essentially Boolean         */
enum atag  ena;    /* enumerated type              */
char       chb;    /* character                    */

float32_t  f32a;   /* 32-bit floating-point - real */
float64_t  f64b;   /* 64-bit floating-point - real */
float128_t f128c;  /* 128-bit floating-point - real */
cfloat32_t cf32a;  /* 32-bit floating-point - complex */
cfloat64_t cf64b; /* 64-bit floating-point - complex */
cfloat128_t cf128c; /* 128-bit floating-point - complex */
```

6.10 Understanding the source references

Where a guideline originates from one or more published sources these are indicated in square brackets after the guideline. This serves two purposes. Firstly the specific sources may be consulted by a reader wishing to gain a fuller understanding of the rationale behind the guideline (for example when considering a request for a deviation). Secondly, with regard to *portability issues* described in the C Standard, the form of the source gives extra information about the nature of the problem.

Rules which do not have a source reference may have originated from a contributing company's in-house standard, or have been suggested by a reviewer, or be widely accepted good practice.

6.10.1 Portability issues

Portability issues are described in subsections of Annexes in the C Standard. It is important to note that the numbering of references is derived from the original C Standard and not from any subsequent revision that incorporates corrections and amendments.

The subsections of the C Standard corresponding to these types of reference are:

Reference	C90	C99, C11
Unspecified	G.1	J.1
Undefined	G.2	J.2
Implementation	G.3	J.3
Locale	G.4	J.4

Where text follows the reference, it has the following meanings:

- Annex G (C90) references: the text identifies the number of an item or items in the relevant section of the Annex, numbered from the beginning of that section. So for example [Locale 2] refers to the second item in section G.4 of the C Standard and [Undefined 3, 5] refers to the third and fifth items in Section G.2 of the C Standard.
- Annex J (C99 and C11) references: for sections J.1, J.2 and J.4, the same interpretation as above applies. For section J.3, the text identifies the subsection of J.3, followed by a parenthesized item number or numbers. So for example [Unspecified 6] refers to the sixth item in Section J.1 of the C Standard and [Implementation J3.4(2, 5)] refers to the second and fifth items in Section J.3.4 of the C Standard.

Where an asterisk (*) is shown, the behaviour is mentioned in the main body of the C Standard, but it is not listed in the appropriate Annex.

Where a guideline is based on issues from Annex G (C90) or Annex J (C99 and C11) of the C Standard, it is helpful for the reader to understand the distinction between the Unspecified, Undefined, Implementation and Locale references. These are explained briefly here, and further information can be found in Hatton [46].

6.10.1.1 Unspecified

These are language constructs that must compile successfully, but in which the compiler writer has some freedom as to what the construct does. An example of this is the “order of evaluation” described in Rule 13.2.

While the use of some unspecified behaviour is unavoidable, it is unwise to assume that the compiler produces object code that behaves in a particular way; the compiler need not even perform consistently across all possible constructs.

MISRA C identifies specific instances of unspecified behaviour which are liable to result in unsafe behaviour.

6.10.1.2 Undefined

These are essentially programming errors, but for which the compiler writer is not obliged to provide error messages. Examples are invalid parameters to functions, or functions whose arguments do not match the defined parameters.

These are particularly important from a safety point of view, as they represent programming errors which may not necessarily be trapped by the compiler.

6.10.1.3 Implementation

These are similar to the “unspecified” issues, the main difference being that the compiler writer must take a consistent approach and document it. In other words the functionality can vary from one compiler to another, making code non-portable, but on any one compiler the behaviour should be well-defined. An example of this is the behaviour of the integer division and remainder operators, / and %, when applied to one positive and one negative integer.

The implementation-defined behaviours tend to be less critical from a safety point of view, provided the compiler writer has fully documented the intended approach and then implemented it consistently. It is advisable to avoid these issues where possible.

6.10.1.4 Locale

The locale-specific behaviours form a small set of features which may vary with international requirements. An example of this is the facility to represent a decimal point by the “,” character instead of the “.” character. No issues arising from this source are addressed in this document.

6.10.2 Other references

References other than from the C Standard *portability issues* are taken from the following sources:

Reference	Source	Ref
MISRA Guidelines	The MISRA Guidelines	[16]
Koenig	C Traps and Pitfalls, Koenig	[48]
IEC 61508	IEC 61508:2010	[36]
ISO 26262	ISO 26262:2018	[34]
DO-178C	DO-178C	[39]
C Secure	ISO/IEC 17961:2013	[33]

Unless stated otherwise, the text following a reference gives the relevant page number.

7 Directives

7.1 The implementation

Dir 1.1 Any implementation-defined behaviour on which the output of the program depends shall be documented and understood

C90 [Annex G.3], C99 [Annex J.3], C11 [Annex J.3]

Category Required

Applies to C90, C99, C11

Amplification

Appendix G of this document lists those implementation-defined behaviours that:

- Are considered to have the potential to cause unexpected program operation, and
- May be present in a program even if it complies with all the other MISRA C guidelines.

All of these implementation-defined behaviours on which the program output depends must be:

- Documented, and
- Understood by developers.

Note: a conforming implementation is required to document its treatment of all implementation-defined behaviour. The developer of an implementation should be consulted if any documentation is missing.

Rationale

It is important to know that the output of a program was intentional and was not produced by chance.

Some of the more common implementation-defined behaviours on which safety-related embedded software is likely to depend are described below.

Core behaviour

The fundamental implementation-defined behaviours, which are likely to be required by most programs, include:

- How to identify a diagnostic message produced during translation;
- The type of the function *main*, commonly declared as `void main (void)` in freestanding implementations;
- The number of significant characters in identifiers — needed to configure an analysis tool for Rule 5.2;
- The source and execution character sets;
- The sizes of the integer types;
- How a *#include'd* name is mapped onto a file name and located in the host file system.

Extensions

Extensions are often used in embedded systems to provide access to peripherals and to place objects into regions of memory with special properties such as Flash EEPROM or fast-access RAM. A conforming implementation is permitted to provide extensions provided that they do not alter the meaning of any strictly conforming program.

An implementation may provide extensions which implement a subset of the features for a later edition of the C Standard.

Some of the methods by which an implementation can provide extensions are:

- The `#pragma` preprocessing directive or the `_Pragma` operator;
- New keywords.

The Standard Library

Some aspects of the Standard Library implementation that may be important are:

- The values assigned to `errno` when certain Standard Library functions are used;
- The implementation of clock and time functions;
- The characteristics of the file system.

The Application Binary Interface

It is sometimes necessary to interface C code with assembly language, for example to improve execution speed in those places where it is critical. It might also be necessary to interface code produced by different compilers, possibly for different languages.

The Application Binary Interface (ABI) for a compiler provides the information necessary to perform this task, including some of the implementation-defined behaviours. It typically specifies:

- How function parameters are passed in registers and on the stack;
- How function values are returned;
- Which registers must be preserved by a function;
- How objects with automatic storage duration are allocated to stack frames;
- Alignment requirements for each data type;
- How structures are laid out and how bit-fields are allocated to storage units.

Some processors have a standard ABI which will be used by all implementations. Where no standard ABI exists, an implementation will provide its own.

Integer division

In C90, signed integer division or remainder operations in which either operand has a negative value may either round downwards or towards zero. In C99 and later, rounding is guaranteed to be towards zero.

Floating-point implementation

The implementation of floating-point types can have significant impact on program behaviour, for example:

- The range of values and the precision with which they are stored;
- The direction of rounding following floating-point operations;
- The direction of rounding when converting to narrower floating-point types or to integer types;
- The behaviour in the presence of underflow, overflow and NaNs;
- The behaviour of library functions in the event of domain and range errors.

See also

Rule 5.1, Rule 5.2

7.2 Compilation and build

Dir 2.1 All source files shall compile without any compilation errors

Category Required

Applies to C90, C99, C11

Rationale

A conforming compiler is permitted to produce an object module despite the presence of compilation errors. However, execution of the resulting program may produce unexpected behaviour.

See also

Rule 1.1

7.3 Requirements traceability

Dir 3.1 All code shall be traceable to documented requirements

[DO-178C Section 6.4.4.3.d]

Category Required

Applies to C90, C99, C11

Rationale

Functionality that is not needed to meet the project requirements gives rise to unnecessary paths. It is possible that the developers of the software are not aware of the wider implications that might arise from this additional functionality. For example, developers might add code that toggles the state of a processor output pin every time a particular point in the program is reached.

This might be very useful during development for measuring timing or for triggering an emulator or logic analyser. However, even though the pin in question might appear to be unused because it is not mentioned in the software requirements specification, it might be connected to an actuator in the target controller, resulting in unwanted external effects.

The method by which code is traced back to documented requirements shall be determined by the project. One method of achieving traceability is to review the code against the corresponding design documents which have in turn been reviewed against the requirements.

Note: there should not be any conflict between this guideline and the provision of a protective coding strategy as the latter should be part of the requirements in a critical system.

7.4 Code design

Dir 4.1 Run-time failures shall be minimized

C90 [Undefined 15, 19, 26, 30, 31, 32, 94]

C99 [Undefined 15, 16, 33, 40, 43–45, 48, 49, 113]

C11 [Undefined 17, 18, 36, 43, 46–48, 51, 52, 119]

Category Required

Applies to C90, C99, C11

Rationale

The C language was designed to provide very limited built-in run-time checking. While this approach allows generation of compact and fast executable code, it places the burden of run-time checking on the programmer. In order to achieve the desired level of robustness, it is therefore important that programmers carefully consider adding dynamic checks wherever there is potential for run-time errors to occur.

It is sometimes possible to demonstrate that the values of operands preclude the possibility of a run-time error during evaluation of an expression. In such cases, a dynamic check is not required provided that the argument supporting its omission is documented. Any such documentation should include the assumptions on which the argument depends. This information can be used during subsequent modifications of the code to ensure that the argument remains valid.

The techniques that will be employed to minimize run-time failures should be planned and documented, for example in design standards, test plans, static analysis configuration files and code review checklists. The nature of these techniques may depend on the integrity requirements of the project.

Note: the presence of a run-time error indicates a violation of Rule 1.3.

The following notes give some guidance on areas where consideration needs to be given to the provision of dynamic checks.

Arithmetic errors

This includes errors occurring in the evaluation of expressions, such as overflow, underflow, divide by zero or loss of significant bits through shifting. In considering integer overflow, note that unsigned integer calculations do not strictly overflow but wrap around producing defined, but possibly unexpected, values.

Careful consideration should be given to the ranges of values and order of operation in arithmetic expressions, for example:

```
float32_t f1 = 1E38f;
float32_t f2 = 10.0f;
float32_t f3 = 0.1f;
float32_t f4 = ( f1 * f2 ) * f3; /* (f1 * f2) will overflow */
float32_t f5 = f1 * ( f2 * f3 ); /* no overflow because (f2 * f3)
                                * is (approximately) 1 */

if ( ( f3 >= 0.0f ) && ( f3 <= 1.0f ) )
{
    /*
     * no overflow because f3 is known to be in range 0..1 so the
     * result of the multiplication will fit in type float32_t
     */
    f4 = f3 * 100.0f;
}
```

Array bound errors

Ensure that array indices are within the bounds of the array size before using them to index the array (Rule 18.1).

Dynamic memory

If dynamic memory allocation is being performed, it is essential to check that each allocation succeeds and that an appropriate degradation or recovery strategy is designed and tested.

Function parameters

The validity of arguments should be checked prior to passing them to library functions (Dir 4.11).

Pointer arithmetic

Ensure that when an address is calculated dynamically the computed address is reasonable and points somewhere meaningful. In particular it should be ensured that if a pointer points within an array, then when the pointer has been incremented or otherwise altered it still points within the same array. See restrictions on pointer arithmetic (Rule 18.1, Rule 18.2 and Rule 18.3).

Pointer dereferencing

Unless a pointer is already known to be non-NULL, a run-time check should be made before dereferencing that pointer. Once a check has been made, it is relatively straightforward within a single function to reason about whether the pointer may have changed and whether another check is therefore required. It is much more difficult to reason across function boundaries, especially when calling functions defined in other source files or libraries.

```
/*
 * Given a pointer to a message, check the message header and return
 * a pointer to the body of the message or NULL if the message is
 * invalid.
 */
const char *msg_body ( const char *msg )
{
    const char *body = NULL;

    if ( msg != NULL )
    {
        if ( msg_header_valid ( msg ) )
        {
            body = &msg[ MSG_HEADER_SIZE ];
        }
    }
    return body;
}
```

```

    char  msg_buffer[ MAX_MSG_SIZE ];
    const char *payload;

    payload = msg_body ( msg_buffer );

    /* Check if there is a payload */
    if ( payload != NULL )
    {
        /* Process the payload */
    }

```

See also

Dir 4.11, Dir 4.12, Rule 1.3, Rule 18.1, Rule 18.2, Rule 18.3

Dir 4.2 All usage of assembly language should be documented

Category Advisory

Applies to C90, C99, C11

Amplification

The rationale for the use of the assembly language and the mechanism for interfacing between C and the assembly language should be documented.

Rationale

Assembly language code is implementation-defined and therefore not portable.

Dir 4.3 Assembly language shall be encapsulated and isolated

Category Required

Applies to C90, C99, C11

Amplification

Where assembly language instructions are used they shall be encapsulated and isolated in:

- Assembly language functions;
- C functions (*inline functions* preferred for C99 and later);
- C macros.

Note: the use of in-line assembly language is an extension to standard C, and therefore violates Rule 1.2.

Rationale

For reasons of efficiency it is sometimes necessary to embed simple assembly language instructions in-line, for example to enable and disable interrupts. If this is necessary, then it is recommended that it be achieved by using macros or *inline functions*.

Encapsulating assembly language is beneficial because:

- It improves readability;
- The name, and documentation, of the encapsulating macro or function makes the intent of the assembly language clear;
- All uses of assembly language for a given purpose can share the same encapsulation which improves maintainability;
- The assembly language can easily be substituted for a different target or for purposes of static analysis.

Example

```
#define NOP asm("    NOP")
```

Dir 4.4 Sections of code should not be “commented out”

Category Advisory

Applies to C90, C99, C11

Amplification

This rule applies to both `//` and `/* ... */` styles of comment.

Rationale

Where it is required for sections of source code not to be compiled then this should be achieved by use of conditional compilation (e.g. `#if` or `#ifdef` constructs with a comment). Using start and end comment markers for this purpose is dangerous because C does not support nested comments, and any comments already existing in the section of code would change the effect.

See also

Rule 3.1, Rule 3.2

Dir 4.5 Identifiers in the same name space with overlapping visibility should be typographically unambiguous

Category Advisory

Applies to C90, C99, C11

Amplification

The definition of the term “unambiguous” should be determined for a project taking into account the alphabet and language in which the source code is being written.

For the Latin alphabet as used in English words, it is advised as a minimum that identifiers should not differ by any combination of:

- The interchange of a lowercase character with its uppercase equivalent;
- The presence or absence of the underscore character;

- The interchange of the letter "O", and the digit "0";
- The interchange of the letter "l", and the digit "1";
- The interchange of the letter "l", and the letter "l" (el);
- The interchange of the letter "l" (el), and the digit "1";
- The interchange of the letter "S", and the digit "5";
- The interchange of the letter "Z", and the digit "2";
- The interchange of the letter "n", and the letter "h";
- The interchange of the letter "B", and the digit "8";
- The interchange of the letter sequence "rn" ("r" followed by "n"), and the letter "m";

Rationale

Depending upon the font used to display the character set, it is possible for certain glyphs to appear the same, even though the characters are different. This may lead to the developer confusing an identifier with another one.

Example

The following examples assume the interpretation suggested in the Amplification for the Latin alphabet and English language.

```
int32_t  id1_a_b_c;
int32_t  id1_abc;      /* Non-compliant */

int32_t  id2_abc;
int32_t  id2_ABC;     /* Non-compliant */

int32_t  id3_a_bc;
int32_t  id3_ab_c;   /* Non-compliant */

int32_t  id4_I;
int32_t  id4_1;     /* Non-compliant */

int32_t  id5_Z;
int32_t  id5_2;     /* Non-compliant */

int32_t  id6_O;
int32_t  id6_0;     /* Non-compliant */

int32_t  id7_B;
int32_t  id7_8;     /* Non-compliant */

int32_t  id8_rn;
int32_t  id8_m;     /* Non-compliant */

int32_t  id9_rn;
struct
{
    int32_t id9_m;      /* Compliant */
};
```

Dir 4.6 *typedefs* that indicate size and signedness should be used in place of the basic numerical types

Category Advisory

Applies to C90, C99, C11

Amplification

The basic numerical types of *char*, *short*, *int*, *long*, *long long*, *float*, *double* and *long double* should not be used, but specific-length *typedefs* should be used. The numerical types of *char* are *signed char* and *unsigned char*. These Guidelines do not treat “plain” *char* as a numerical type (see Section 8.10.2 on *essentially character* types).

For C99 and later, the types provided by `<stdint.h>` should be used. For C90, equivalent types should be defined and used.

A type must not be defined with a specific length unless the implemented type is actually of that length.

It is not necessary to use *typedefs* in the declaration of bit-fields.

For example, on a C90 implementation (where the standard header `<stdint.h>` is not available) the following definitions may be suitable, but need to be adjusted to suit the project’s implementation:

```
typedef signed   char    int8_t;
typedef signed  short   int16_t;
typedef signed   int    int32_t;
typedef signed   long   int64_t;

typedef unsigned char   uint8_t;
typedef unsigned short  uint16_t;
typedef unsigned int    uint32_t;
typedef unsigned long   uint64_t;

typedef float         float32_t;
typedef double        float64_t;
typedef long double   float128_t;

typedef float  _Complex cfloat32_t;
typedef double _Complex cfloat64_t;
typedef long double _Complex cfloat128_t;
```

Rationale

In situations where the amount of memory being allocated is important, using specific-length types makes it clear how much storage is being reserved for each object.

Adherence to this guideline does **not** guarantee portability because the size of the *int* type may determine whether or not an expression is subject to integer promotion. For example, an expression with type `int16_t` will not be promoted if *int* is implemented using 16 bits but will be promoted if *int* is implemented using 32 bits. This is discussed in more detail in the section on integer promotion in Appendix C.

Note: defining a specific-length type whose size is **not** the same as the implemented type is counter-productive both in terms of storage requirements and in terms of portability. Care should be taken to avoid defining types with the wrong size.

If abstract types are defined in terms of a specific-length type then it is not necessary, and may even be undesirable, for those abstract types to specify the size or sign. For example, the following code defines an abstract type representing mass in kilograms but does not indicate its size or sign:

```
typedef uint16_t mass_kg_t;
```

It might be desirable not to apply this guideline when interfacing with the Standard Library or code outside the project's control.

Exception

1. The basic numerical types may be used in a *typedef* to define a specific-length type.
2. For function `main` an *int* may be used rather than the *typedefs* as a return type. Therefore `int main (void)` is permitted.
3. For function `main` an *int* may be used rather than the *typedefs* for the input parameter `argc`.

Therefore `int main(int argc, char *argv[])` is permitted.

Example

```
/* Non-compliant - int used to define an object */
int x = 0;

/* Compliant - int used to define specific-length type */
typedef int SINT_16;

/* Non-compliant - no sign or size specified */
typedef int speed_t;

/* Compliant - further abstraction does not need specific length */
typedef int16_t torque_t;
```

Dir 4.7 If a function returns error information, then that error information shall be tested

Category Required

Applies to C90, C99, C11

Amplification

The list of functions that are deemed to return error information shall be determined by the project.

The error information returned by a function shall be tested in a meaningful manner.

Rationale

A function (whether it is part of the Standard Library, a third party library or a user defined function) may be deemed to provide some means of indicating the occurrence of an error. This may be via an error flag, some special return value or some other means. Whenever such a mechanism is provided by a function the calling program shall check for the indication of an error as soon as the function returns.

However, note that the checking of input values to functions is considered a more robust means of error prevention than trying to detect errors after the function has completed (see Dir 4.11).

Exception

If it can be shown, for example by checking arguments, that a function cannot return an error indication then there is no need to perform a check.

See also

Dir 4.11, Rule 17.7

Dir 4.8 If a pointer to a structure or union is never dereferenced within a translation unit, then the implementation of the object should be hidden

Category Advisory

Applies to C90, C99, C11

Amplification

The implementation of an object should be hidden by means of a pointer to an incomplete type. This directive only applies if all the pointers to a particular structure or union in a translation unit are never dereferenced.

Rationale

If a pointer to a structure or union is never dereferenced, then the implementation details of the object are not needed and its contents should be protected from unintentional changes.

Hiding the implementation details creates an *opaque type* which may be referenced via a pointer but whose contents may not be accessed.

Example

```
/* Opaque.h */

#ifndef OPAQUE_H
#define OPAQUE_H

typedef struct OpaqueType *pOpaqueType;

#endif

/* Opaque.c */

#include "Opaque.h"

struct OpaqueType
{
    /* Object implementation */
};

/* UseOpaque.c */

#include "Opaque.h"

void f ( void )
{
    pOpaqueType pObj;

    pObj = GetObject ( ); /* Get a handle to an OpaqueType object */

    UseObject ( pObj ); /* Use it... */
}
```

Dir 4.9 A function should be used in preference to a *function-like macro* where they are interchangeable

[Koenig 78–81]

Category Advisory

Applies to C90, C99, C11

Amplification

This guideline applies only where a function is permitted by the syntax and *constraints* of the language standard.

Rationale

In most circumstances, functions should be used instead of macros. Functions perform argument type-checking and evaluate their arguments once, thus avoiding problems with potential multiple *side effects*. In many debugging systems, it is easier to step through execution of a function than a macro. Nonetheless, macros may be useful in some circumstances.

Some of the factors that should be considered when deciding whether to use a function or a macro are:

- The benefits of function argument and result type-checking;
- The availability of *inline functions*, although note that the extent to which *inline* is acted upon is implementation-defined;
- The trade-off between code size and execution speed;
- Whether the possibility for compile-time evaluation is important: macros with constant arguments are more likely to be evaluated at compile-time than corresponding function calls;
- Whether the arguments would be valid for a function: macro arguments are textual whereas function arguments are expressions;
- Ease of understanding and maintainability.

Exception

The use of *function-like macros* for `_Generic` selections is advised by Rule 23.1, and is therefore permitted by this Directive.

Example

The following example is compliant. The *function-like macro* cannot be replaced with a function because it has a C operator as an argument:

```
#define EVAL_BINOP( OP, L, R ) ( ( L ) OP ( R ) )
uint32_t x = EVAL_BINOP ( +, 1, 2 );
```

In the following example, the use of a macro to initialize an object with static storage duration is compliant because a function call is not permitted here:

```
#define DIV2(X) ( ( X ) / 2 )

void f ( void )
{
    static uint16_t x = DIV2 ( 10 ); /* Compliant - call not permitted */
    uint16_t y = DIV2 ( 10 ); /* Non-compliant - call permitted */
}
```

The following is compliant by exception – in most cases `_Generic` selections cannot be replaced with a function.

```
#define GFUNC(X) ( _Generic( ... ) ) /* Compliant by exception */
```

See also

Rule 13.2, Rule 20.7, Rule 23.1

Dir 4.10 Precautions shall be taken in order to prevent the contents of a *header file* being included more than once

Category Required

Applies to C90, C99, C11

Rationale

When a translation unit contains a complex hierarchy of nested *header files*, it is possible for a particular *header file* to be included more than once. This can be, at best, a source of confusion. If this multiple inclusion leads to multiple or conflicting definitions, then this can result in undefined or erroneous behaviour.

Example

```
/* file.h */
#ifndef FILE_H
/* Non-compliant - does not #define FILE_H */
#endif
```

The following examples show two ways by which the contents of a header file could be protected from being included more than once in a translation unit, but this is not an exclusive list.

```
<start-of-file>
#if !defined ( identifier )
#define identifier
    /* Contents of file */
#endif
<end-of-file>
```

```
<start-of-file>
#ifndef identifier
#define identifier
    /* Contents of file */
#endif
<end-of-file>
```

Notes:

1. The identifier used to test and record whether a given *header file* has already been included shall be unique across all *header files* in the project.
2. Comments are permitted anywhere within these forms.

Dir 4.11 The validity of values passed to library functions shall be checked

C90 [Undefined 60, 63, 96; Implementation 45–47]
 C99 [Unspecified 30, 31, 44, 48–50;
 Undefined 102, 103, 107, 112, 180, 181, 183, 187, 189;
 Implementation J.3.12(8–11)]
 C11 [Unspecified 30, 31, 44, 52–56;
 Undefined 108, 109, 113, 118, 191, 192, 194, 199, 1201;
 Implementation J.3.12(8–14)]

Category Required

Applies to C90, C99, C11

Amplification

The nature and organization of the project will determine which libraries, and functions within those libraries, should be subject to this directive.

Rationale

Many functions in the Standard Library are not required by the C Standard to check the validity of parameters passed to them. Even where checking is required by the C Standard, or where compiler writers claim to check parameters, there is no guarantee that adequate checking will take place.

Similarly, the interface description for functions in other libraries may not specify the checks performed by those functions. There is also a risk that the specified checks are not performed adequately.

The programmer shall provide appropriate checks of input values for all library functions which have a restricted input domain (the Standard Library, third-party libraries, and in-house libraries).

Examples of functions from the Standard Library that have a restricted domain and need checking are:

- Many of the maths functions in `<math.h>`, for example:
 - Negative numbers must not be passed to the *sqrt* or *log* functions;
 - The second parameter of *fmod* should not be zero;
- Some implementations can produce unexpected results when the function *toupper* is passed an argument which is not a lowercase letter (and similarly for *tolower*);
- The character testing functions in `<ctype.h>` exhibit undefined behaviour if passed invalid values;
- The *abs* function applied to the most negative integer gives undefined behaviour.

Although most of the math library functions in `<math.h>` define allowed input domains, the values they return when a domain error occurs may vary from one compiler to another. Therefore pre-checking the validity of the input values is particularly important for these functions.

The trigonometric periodic functions in `<math.h>` incur significant precision loss when called with arguments with relatively large absolute value. If x is an IEC 60559 single-precision number and $x \geq 2^{23}$, then the smallest single-precision range containing $[x, x + 2\pi)$ contains no more than three floating-point numbers (the same holds for IEC 60559 [35] double-precision numbers substituting 2^{52} to 2^{23}). This implies that computing trigonometric periodic functions on large values, because of the gross input inaccuracy, gives non-significant results independently from the quality of the function implementation. For this reason, trigonometric periodic functions should not be called on arguments whose absolute value is larger than $k\pi$, where k is a property of the floating-point representation. Ideally, this principle should be applied with $k = 1$, but depending on the application and its precision goals, a larger value for k can be used. The programmer should identify any domain constraints which should sensibly apply to a function being used (which may or may not be documented in the interface description), and provide appropriate checks that the input value(s) lies within this domain. Of course the value may be restricted further, if required, by knowledge of what the parameter represents and what constitutes a sensible range of values for the parameter.

There are a number of ways in which the requirements of this guideline might be satisfied, including the following:

- Check the values before calling the function;
- Check the values in the called library function — this is particularly applicable for in-house designed libraries, though it could apply to bought-in libraries if the supplier can demonstrate that they have built in the checks;
- Produce “wrapped” versions of functions that perform the checks then call the original function;
- Demonstrate statically that the input parameters can never take invalid values.

See also

Dir 4.1, Dir 4.7

Dir 4.12 Dynamic memory allocation shall not be used

Category Required

Applies to C90, C99, C11

Amplification

This rule applies to all dynamic memory allocation packages including:

- Those provided by the Standard Library;
- Third-party packages.

Rationale

The Standard Library's dynamic memory allocation and deallocation routines can lead to undefined behaviour as described in Rule 21.3. Any other dynamic memory allocation system is likely to exhibit undefined behaviours that are similar to those of the Standard Library.

The specification of third-party routines shall be checked to ensure that dynamic memory allocation is not being used inadvertently.

If a decision is made to use dynamic memory, care shall be taken to ensure that the software behaves in a predictable manner. For example, there is a risk that:

- Insufficient memory may be available to satisfy a request — care must be taken to ensure that there is a safe and appropriate response to an allocation failure;
- There is a high variance in the execution time required to perform allocation or deallocation depending on the pattern of usage and resulting degree of fragmentation.

Example

For convenience, these examples are based around use of the Standard Library's dynamic memory functions as their interfaces are well-known.

In this example, the behaviour is undefined following the first call to *free* because the value of the pointer *p* becomes indeterminate. Although the value stored in the pointer is unchanged following the call to *free*, it is possible, on some targets, that the memory to which it points no longer exists and the act of copying that pointer could cause a memory exception.

```
#include <stdlib.h>

void f ( void )
{
    char    *p = ( char * ) malloc ( 10 );
    char    *q;

    free ( p );
    q = p;      /* Undefined behaviour - value of p is indeterminate */

    p = ( char * ) malloc ( 20 );
    free ( p );
    p = NULL;   /* Assigning NULL to freed pointer makes it determinate */
}
```

See also

Dir 4.1, Rule 18.7, Rule 21.3, Rule 22.1, Rule 22.2

Dir 4.13 Functions which are designed to provide operations on a resource should be called in an appropriate sequence

Category Advisory

Applies to C90, C99, C11

Amplification

A set of functions providing operations on a resource typically has three kinds of operation:

1. Allocation of the resource, e.g. opening a file;
2. Deallocation of the resource, e.g. closing a file;
3. Other operations, e.g. reading from a file.

For each such set of functions, all uses of its operations should occur in an appropriate sequence.

Rationale

Static analyser tools are capable of providing path analysis checks that can identify paths through a program that result in the deallocation function of a sequence not being called. In order to maximize the benefits of such automated checks, developers are therefore encouraged to enable these checks by designing and declaring sets of balanced functions to the static analyser.

Example

```
/* These functions are intended to be paired */
extern mutex_t mutex_lock ( void );
extern void    mutex_unlock ( mutex_t m );

extern int16_t x;

void f ( void )
{
    mutex_t m = mutex_lock ( );

    if ( x > 0 )
    {
        mutex_unlock ( m );
    }
    else
    {
        /* Mutex not unlocked on this path */
    }
}
```

See also

Rule 22.1, Rule 22.2, Rule 22.6

Dir 4.14 The validity of values received from external sources shall be checked

C90 [Undefined 15, 19, 26, 30, 31, 32, 94]

C99 [Undefined 15, 16, 33, 40, 43–45, 48, 49, 113]

C11 [Undefined 17, 18, 36, 43, 46–48, 51, 52, 119]

Category Required

Applies to C90, C99, C11

Amplification

“External sources” include data:

- Read from a file;
- Read from an environment variable;
- Resulting from user input;
- Received over a communications channel.

Rationale

A program has no control over the values given to data originating from external sources. The values may therefore be invalid, either as the result of errors or due to malicious modification by an external agent. Data from external sources shall therefore be validated before it is used.

In the security domain, external sources of data are usually regarded as untrusted as they may have been modified by someone trying to harm or gain control of the program and/or system it is running on; such data needs to be validated before it can be used safely.

In the safety domain, external sources are regarded as “suspicious” and values obtained from them require validation.

In both domains, data from an external source shall be tested to ensure that its value respects all the constraints placed on its use (i.e. its value is not harmful), even if the value cannot be proven to be correct. For example:

- A value used to compute an array index shall not result in an array bounds error;
- A value used to control a loop shall not cause excessive (e.g. infinite) iteration;
- A value used to compute a divisor shall not result in division by zero;
- A value used to compute an amount of dynamic memory shall not result in excessive memory allocation;
- A string used as a query to an SQL database shall be checked to ensure that it does not include a ; character.

Example

The following example is non-compliant as there is no check made to ensure that a string resulting from user input is null terminated. This may lead to an array bounds error, commonly known as a buffer overrun, when the string is output through the call to *printf*.

```
void f1( void )
{
    char input [ 128 ];
    ( void ) scanf ( "%128c", input );
    ( void ) printf ( "%s", input );      /* Non-compliant */
}
```

Dir 4.15 Evaluation of floating-point expressions shall not lead to the undetected generation of infinities and NaNs

C90 [Implementation 20, 22]

C99 [Unspecified 30, 46, 47, 48, 49; Implementation J.3.6(2), J.3.6(5)]

C11 [Unspecified 30, 52, 53, 56, 57; Implementation J.3.6(2), J.3.6(5)]

Category Required

Applies to C90, C99, C11

Amplification

If the evaluation of an arithmetic floating-point expression or a call to a mathematical function can possibly generate an infinity or a NaN (not-a-number), then that result shall be adequately tested. For efficiency reasons, it is possible to exploit the implementation-defined propagation rules for infinities and NaNs to delay the test, provided no infinity or NaN reaches sections of code that have not been designed to handle them.

Rationale

Operations on floating-point numbers can overflow, that is, generate results of magnitude too large to be represented in the considered format. For such cases, the widely-used IEC 60559 Standard has special representations for positive and negative infinity. While the possibility to denote infinities is useful in some contexts because it allows operations to continue after overflow has occurred, much care has to be exercised, because:

- It is likely that an error has occurred that resulted in an approximation error of unknown magnitude;
- Infinities can easily generate unwanted underflows;
- Careless computation with infinities can cause the generation of NaNs, whose correct handling requires even more care.

Note: an infinity does not necessarily correspond to a *large* value: for instance, on an implementation that conforms to IEC 60559, a float computation whose exact result is slightly larger than $3.4e+38$ will be rounded to infinity and `logf(infinity)` returns infinity even though `logf(3.4e+38)` is less than 89.

Some operations on floating-point numbers are invalid or undetermined, such as taking the square root of a negative number or dividing a zero by a zero. The IEC 60559 Standard has special representations for such exceptional, non-numerical results: they are called NaNs. While of course there are legitimate uses of NaNs, programming with NaNs requires exceptional care, for example:

- In some implementations there are both “signalling” and “quiet” NaNs;
- NaNs can be signed;
- `x == x` is false if `x` is a NaN;
- Some library functions (such as `hypot()`) can return a non-NaN even if one of the arguments is NaN.

For these reasons, projects that do use infinities or NaNs need to take measures ensuring such special representations are detected and handled before they propagate to areas of the systems that are not prepared to receive them.

Example

```
#include <math.h>

extern float64_t get_result();           /* Return value may be infinite
... or valid data */
extern void      use_result( float64_t c ); /* Not protected against NaNs or
... or infinities */

void f( void )
{
    float64_t a = get_result();

    use_result( a );      /* Non-compliant - a not tested */

    if ( !isnan( a ) )
    {
        use_result( a ); /* Non-compliant - a may be infinite */
    }

    if ( isfinite( a ) )
    {
        use_result( a ); /* Compliant */
    }
}
```

The following example shows that it is not necessary to check for NaNs or infinities at every step:

```
void g( void )
{
    float64_t a = get_result();

    float64_t b = exp( -2.0 * a ); /* Compliant - exp() propagates infinities
... and NaNs as expected */

    float64_t c = ( a * ( 1.0 + b ) - ( 1.0 - b ) ) / ( a * a );

    /* Division can result in NaN, even if operands are not infinity or NaN */

    use_result( c ); /* Non-compliant - not protected against NaNs
... or infinities */

    if ( isfinite( c ) )
    {
        use_result( c ); /* Compliant - protected against NaNs
... and infinities */
    }
}
```

See also

Dir 1.1, Dir 4.1, Dir 4.7

7.5 Concurrency considerations

Dir 5.1 There shall be no data races between threads

C11 [Undefined 5, *]

Category Required

Applies to C11

Amplification

Two expression evaluations conflict if one of them modifies a memory location and the other one reads or modifies the same memory location. The execution of a program contains a data race if it contains two conflicting actions in different threads, at least one of which is not atomic, and neither happens before the other, i.e. there is no fixed ordering between the two actions. To prevent data races, objects shared between different threads shall be protected by an appropriate synchronization mechanism.

Rationale

Data races are caused by simultaneous accesses to the same non-atomic object from two different threads $T1$ and $T2$ where at least one of them is a write access and where the program semantics does not impose a fixed ordering between $T1$ and $T2$. There may be legitimate program executions where the access from $T1$ is executed before the access from thread $T2$, and vice versa, or where a given access itself is interrupted. Any such data race results in undefined behaviour.

There are several critical scenarios:

- Depending on the timing of the threads, sometimes in a given context the wrong value might be used, leading to unexpected results.
- If a read or write access is implemented by several machine instructions, a pre-emption might occur between these instructions such that inconsistent values might be read or written. As an example, a 64-bit variable read implemented as two 32-bit load instructions might be interrupted after reading the first 32 bits. Then another thread might change the variable value. When the first thread resumes, it reads the second 32-bit half, which now contains a different value than when the first 32 bits of the variable were read.

In general, a data race can cause memory corruption and lead to unexpected, erroneous or erratic behaviour. Data races typically manifest sporadically and are very hard to reproduce.

To prevent such situations, when an object is shared between different threads, it shall be protected by an appropriate synchronization mechanism. To ensure consistent access within a single shared object it can be declared as atomic. A more general solution to ensure consistency of accesses is to introduce critical sections with mutex locks or condition variables.

Note: C Standard Library functions may access objects with static or thread storage duration directly or indirectly via the function's arguments. The C Standard Library functions *setlocale*, *tmpnam*, *rand*, *srand*, *getenv*, *getenv_s*, *strtok*, *strerror*, *asctime*, *ctime*, *gmtime*, *localtime*, *mbrtoc16*, *c16rtomb*, *mbrtoc32*, *c32rtomb*, *mbrlen*, *mbrtowc*, *wcrtomb*, *mbsrtowcs*, *wcsrtombs* are not guaranteed to be reentrant and may modify objects with static or thread storage duration. To prevent data races explicit synchronization may be required.

Example

The following example exhibits data races on the global variables `x` and `a`. Functions `t1`, `t2`, `t3` and `t4` are executed as concurrent threads `T1`, `T2`, `T3` and `T4` respectively.

Variable `x` is accessed without synchronization, by function `t1` in thread `T1` and by function `t2` in thread `T2`. If executed on a 16-bit machine writing 32-bit values in two chunks of 16 bits, threads `T1` and `T2` might interrupt one another after the first 16 bits of the variable have been written. As a consequence, the two 16-bit halves of variable `x` might be written by different threads, causing unexpected values.

```
int32_t x;
int32_t a = 1;
int32_t b;

int32_t t1( void *ignore ) /* Thread T1 entry */
{
    while ( 1 )
    {
        x = -1; /* Write-write data race with t2. Possible values of x: 0xFFFF0000,
                0x0000FFFF, 0x00000000, 0xFFFFFFFF */
    }
    return 0;
}

int32_t t2( void *ignore ) /* Thread T2 entry */
{
    while ( 1 )
    {
        x = 0; /* Write-write data race with t1. Possible values of x: 0xFFFF0000,
                0x0000FFFF, 0x00000000, 0xFFFFFFFF */
    }
    return 0;
}
```

A data race on `a` is caused by unprotected accesses by function `t3` in thread `T3` and by function `t4` in thread `T4`. If thread `T3` sees the value of 1 in variable `a`, it will enter the then-part of the conditional statement. At that point, it might be interrupted by thread `T4`, which sets `a` to 0. After resuming, thread `T3` will run into a division by zero.

```
int32_t t3( void *ignore ) /* Thread T3 entry */
{
    while ( 1 )
    {
        if ( a != 0 ) /* Read-write data race with T4 */
        {
            b += 1/a; /* Read-write data race with T4 */
            a = 1; /* Write-write data race with T4 */
        }
    }
    return 0;
}

int32_t t4( void *ignore ) /* Thread T4 entry */
{
    while ( 1 )
    {
        a = 0; /* Read-write data race with T3 */
    }
    return 0;
}
```

See also

Rule 9.7, Rule 12.6

Category Required

Applies to C11

Amplification

A deadlock occurs when there is a circular chain of threads, each of which holding a locked synchronization resource, and trying to lock a synchronization resource held by the next element in the chain. To prevent deadlocks, synchronization mechanisms between threads shall not introduce cyclic dependencies.

Rationale

An example for a deadlock between two threads $T1$ and $T2$ is when $T1$ enters the waiting state because it requests a mutex R_a which is locked by thread $T2$, and $T2$ in turn is waiting for another mutex R_b held by thread $T1$.

Possible solutions to avoid deadlocks include locking and unlocking synchronization resources in a fixed global non-cyclic order, or associating synchronization resources with appropriate priorities.

Example

Assume that in the following example functions $t1$ and $t2$ are executed as concurrent threads $T1$ and $T2$. Thread $T1$ locks mutex R_a , then executes some other code in which it might be interrupted by thread $T2$. Thread $T2$ locks mutex R_b , executes some other code, and is blocked when attempting to lock mutex R_a , which is currently held by thread $T1$. Hence thread $T1$ resumes, and eventually reaches the call to `mtx_lock(&Rb)` on which it blocks, because R_b is held by $T2$. Then execution is stuck indefinitely because thread $T1$ is waiting for thread $T2$ and vice versa.

```

mtx_t  Ra;
mtx_t  Rb;

int32_t t1( void *ignore ) /* Thread T1 entry      */
{
    mtx_lock( &Ra );
    ...
    mtx_lock( &Rb );          /* Deadlock may occur here */
    ...
    mtx_unlock( &Rb );
    mtx_unlock( &Ra );
    return 0;
}

int32_t t2( void *ignore ) /* Thread T2 entry      */
{
    mtx_lock( &Rb );
    ...
    mtx_lock( &Ra );          /* Deadlock may occur here */
    ...
    mtx_unlock( &Ra );
    mtx_unlock( &Rb );
    return 0;
}

```

Dir 5.3 There shall be no dynamic thread creation

Category Required

Applies to C11

Amplification

Thread creation shall only occur in a well-defined program start-up phase.

Rationale

Uncertainty about the number of threads running at a particular point in time is error prone and reduces analysability. Also the overhead in thread creation and destruction is hard to predict.

Usage of a static thread pool is common practice in operating systems for safety-related systems, e.g. ARINC-653 [38], AUTOSAR [42] and OSEK [49].

Example

```

thrd_t id1;
thrd_t id2;

int32_t t1( void *ignore )      /* Thread T1 entry          */
{
    ...
    thrd_create( &id2, t2, NULL ); /* Non-compliant, not constrained to start-up */
    ...
}

int32_t t2( void *ignore )      /* Thread T2 entry          */
{
    ...
}

void main(void)
{
    thrd_create( &id1, t1, NULL ); /* Compliant                */
    ...
}

```

See also

Dir 4.7

8 Rules

8.1 A standard C environment

Rule 1.1 The program shall contain no violations of the standard C syntax and *constraints*, and shall not exceed the implementation's translation limits

[MISRA Guidelines Table 3], [IEC 61508-7: Table C.1], [ISO 26262-6: Table 1]

Category Required
Analysis Decidable, Single Translation Unit
Applies to C90, C99, C11

Amplification

The program shall use only those features of the C Standard and its Standard Library that are specified in the chosen edition of the C Standard (see Section 1.3).

The C Standard permits implementations to provide language extensions and the use of such extensions is permitted by this rule.

Except when making use of a language extension, a program shall not:

- Contain any violations of the language syntax described in the C Standard;
- Contain any violations of the *constraints* imposed by the C Standard.

A program shall not exceed the translation limits imposed by the implementation. The minimum translation limits are specified by the C Standard but an implementation may provide higher limits.

Note: a conforming implementation generates a diagnostic for syntax and *constraint* violations but be aware that:

- The diagnostic need not necessarily be an error but could, for example, be a warning;
- The program may be translated and an executable generated despite the presence of a syntax or *constraint* violation;

Note: a conforming implementation does not need to generate a diagnostic when a translation limit is exceeded; an executable may be generated but it is not guaranteed to execute correctly.

Rationale

Problems associated with language features that are outside the supported editions of the C Standard have not been considered during development of these guidelines.

There is anecdotal evidence of some non-conforming implementations failing to diagnose *constraint* violations, for example in [50] p135, example 2 entitled "Error of writing into the const area".

Example

Some C90 compilers provide support for *inline functions* using the `__inline` keyword. A C90 program that uses `__inline` will be compliant with this rule provided that it is intended to be translated using such a compiler.

Many compilers for embedded targets provide additional keywords that qualify object types with attributes of the memory area in which the object is located, for example:

- `__zpage` — the object can be accessed using a short instruction
- `__near` — a pointer to the object can be held in 16 bits
- `__far` — a pointer to the object can be held in 24 bits

A program using these additional keywords will be compliant with this rule provided that the compiler supports those keywords as a language extension.

See also

Dir 2.1, Rule 1.2

Rule 1.2 Language extensions should not be used

Category Advisory

Analysis Undecidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

A program that relies on language extensions may be less portable than one that does not. Although the C Standard requires that a conforming implementation document any extensions that it provides to the language, there is a risk that this documentation might not provide a full description of the behaviour in all circumstances.

If this rule is not applied, the decision to use each language extension should be justified in the project's design documentation. The methods by which valid use of each extension will be assured, for example checking the compiler and its diagnostics, should also be documented.

It is recognized that it is necessary to use language extensions in embedded systems. The C Standard requires that an extension does not alter the behaviour of any strictly conforming program. For example, a compiler might implement, as an extension, full evaluation of binary logical operators even though the C Standard specifies that evaluation stops as soon as the result can be determined. Such an extension does not conform to the C Standard because *side effects* in the right-hand operand of a logical AND operator would always occur, giving rise to a different behaviour.

See also

Rule 1.1

Rule 1.3 There shall be no occurrence of undefined or critical unspecified behaviour

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

Some undefined and unspecified behaviours are dealt with by specific rules. This rule prevents all other undefined and critical unspecified behaviours. Appendix H lists the undefined behaviours and those unspecified behaviours that are considered critical.

Rationale

Any program that gives rise to undefined or unspecified behaviour may not behave in the expected manner. In many cases, the effect is to make the program non-portable but it is also possible for more serious problems to occur. For example, undefined behaviour might affect the result of a computation. If correct operation of the software is dependent on this computation then system safety might be compromised. The problem is particularly difficult to detect if the undefined behaviour only manifests itself on rare occasions.

Many of the MISRA C guidelines have been designed to avoid certain undefined and unspecified behaviours. However, other behaviours are not covered by specific guidelines, for example because:

- It is unlikely that the behaviour will be encountered;
- There is no practical guidance that can be given other than the obvious statement that the behaviour should be avoided.

Instead of introducing a guideline for each undefined and critical unspecified behaviour, the MISRA C Guidelines directly address those that are considered most important and most likely to occur in practice. Those behaviours that do not have specific guidelines are all covered together by this single rule. Appendix H lists all undefined and critical unspecified behaviours, along with the MISRA C guidelines that prevent their occurrence. It therefore indicates which behaviours are expected to be prevented by this rule and which behaviours are covered by other rules.

Note: some implementations may provide well-defined behaviour for some of the undefined and unspecified behaviours listed in the C Standard. If such well-defined behaviours are relied upon, including by means of a language extension, it will be necessary to deviate this rule in respect of those behaviours.

See also

Dir 4.1

Rule 1.4 Emergent language features shall not be used

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C11

Amplification

Updates to the C Standard have introduced new language features. The following shall not be used:

- Other than defining `__STDC_WANT_LIB_EXT1__` to 0, the facilities of Annex K (Bounds-checking interfaces) shall not be used.

Rationale

Use of the language features restricted by this rule may have instances of undefined, unspecified or implementation-defined behaviour associated with them. In addition, features may also exhibit well defined behaviour that does not meet developer expectations.

Any instances of undefined, unspecified or implementation-defined behaviour are diagnosed by:

- Dir 1.1, which requires that the use of implementation-defined behaviour be documented and taken into consideration; and
- Rule 1.3, which prohibits the presence of undefined and critical unspecified behaviours.

However, detection of these behaviours alone does not mitigate against well-defined behaviour that does not meet developer expectations. Additional static analysis checks are needed to check for these behaviours, but there is no requirement for a static analysis tool to implement these checks as they are not specified within The Guidelines.

This Rule requires that any use of an emergent language feature be supported by a deviation to ensure that all undesirable behaviours are identified and measures put in place to ensure that they do not compromise safety or security.

See also

Dir 1.1, Rule 1.3

Rule 1.5 Obsolescent language features shall not be used

Category	Required
Analysis	Undecidable, System
Applies to	C99, C11

Amplification

Obsolescent features are those identified in the *Future language directions* and *Future library directions* sections of the C Standard, and are listed in Appendix F.

Rationale

Features are declared as obsolescent by the C Standard when they are superseded by safer or better alternatives, or are considered to exhibit undesirable behaviour. Features declared as obsolescent by a particular edition of the C Standard may be withdrawn in a later edition.

See also

Rule 1.1

8.2 Unused code

Rule 2.1 A project shall not contain *unreachable code*

[IEC 61508-7 Section C.5.9], [DO-178C Section 6.4.4.3.c]

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Rationale

Provided that a program does not exhibit any undefined behaviour, *unreachable code* cannot be executed and cannot have any effect on the program's outputs. The presence of *unreachable code* may therefore indicate an error in the program's logic.

A compiler is permitted to remove any *unreachable code* although it does not have to do so. *Unreachable code* that is **not** removed by the compiler wastes resources, for example:

- It occupies space in the target machine's memory;
- Its presence may cause a compiler to select longer, slower jump instructions when transferring control around the *unreachable code*;
- Within a loop, it might prevent the entire loop from residing in an instruction cache.

It is sometimes desirable to insert code that appears to be unreachable in order to handle exceptional cases. For example, in a *switch* statement in which every possible value of the controlling expression is covered by an explicit *case*, a *default* clause shall be present according to Rule 16.4. The purpose of the *default* clause is to trap a value that should not normally occur but that may have been generated as a result of:

- Undefined behaviour present in the program;
- A failure of the processor hardware.

If a compiler can prove that a *default* clause is unreachable, it may remove it, thereby eliminating the defensive action. On the assumption that the defensive action is important, it will be necessary either to demonstrate that the compiler does not eliminate the code despite it being unreachable, or to take steps to make the defensive code reachable. The former course of action requires a deviation against this rule, probably with a review of the object code or unit testing being used to support such a deviation. The latter course of action can usually be achieved by means of a *volatile* access. For example, a compiler might determine that the range of values held by *x* is covered by the *case* clauses in a *switch* statement such as:

```
uint16_t x;
switch ( x )
```

By forcing *x* to be accessed by means of a *volatile* qualified *lvalue*, the compiler has to assume that the controlling expression could take any value:

```
switch ( *( volatile uint16_t * ) &x )
```

Note: code that has been conditionally excluded by pre-processor directives is not subject to this rule as it is not presented to the later phases of translation.

Example

```
enum light { red, amber, red_amber, green };
enum light next_light ( enum light c )
{
    enum light res;

    switch ( c )
    {
        case red:
            res = red_amber;
            break;
        case red_amber:
            res = green;
            break;
        case green:
            res = amber;
            break;
        case amber:
            res = red;
            break;

        default:
        {
            /*
             * This default will only be reachable if the parameter c
             * holds a value that is not a member of enum light.
             */
            error_handler ( );
            break;
        }
    }

    return res;
    res = c;                                /* Non-compliant - this statement is
                                             * certainly unreachable */
}
```

See also

Rule 14.3, Rule 16.4

Rule 2.2 A project shall not contain *dead code*

[IEC 61508-7 Section C.5.10], [ISO 26262-6 Section 9.4.5], [DO-178C Section 6.4.4.3.c]

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

Any operation that is executed but whose removal would not affect program behaviour constitutes *dead code*. Operations that are introduced by language extensions are assumed always to have an effect on program behaviour.

Notes

1. The behaviour of an embedded system is often determined not just by the nature of its actions, but also by the time at which they occur.
2. *unreachable code* is not *dead code*, as it cannot be executed.
3. Initialization is not the same as an assignment operation and is therefore not a candidate for *dead code*.

Rationale

The presence of *dead code* may be indicative of an error in the program's logic. Since *dead code* may be removed by a compiler, its presence may cause confusion.

Exception

1. A cast to *void* is assumed to indicate a value that is intentionally not being *used*. The cast is therefore not *dead code* itself. It is treated as using its operand which is therefore also not *dead code*.
2. A cast operator whose result is *used* is not *dead code*.

Example

In this example, it is assumed that the object pointed to by *p* is used in other functions.

```
extern volatile uint16_t v;

extern char *p;

void f ( void )
{
    uint16_t x;

    ( void ) v;          /* Compliant      - v is accessed for its side effect
                        *                  and the cast to void is permitted
                        *                  by exception                               */
    ( int32_t ) v;      /* Non-compliant - the cast operator is dead */
    v >> 3;             /* Non-compliant - the >> operator is dead */
    x = 3;              /* Non-compliant - the = operator is dead
                        *                  - x is not subsequently read          */
    *p++;              /* Non-compliant - result of * operator is not used */
    ( *p )++;          /* Compliant      - *p is incremented          */
}
```

In the following compliant example, the `__asm` keyword is a language extension, not a function call operation, and is therefore not *dead code*:

```
__asm ( "NOP" );
```

In the following example, the function `g` does not contain *dead code*, and is not itself *dead code* because it does not contain any operations. However, the call to the function is dead because it could be removed without affecting program behaviour.

```
void g ( void )
{
  /* Compliant - there are no operations in this function */
}

void h ( void )
{
  g ( ); /* Non-compliant - the call could be removed */
}
```

See also

Rule 17.7

Rule 2.3 A project should not contain unused type declarations

Category Advisory

Analysis Decidable, System

Applies to C90, C99, C11

Rationale

If a type is declared but not used, then it is unclear to a reviewer if the type is redundant or it has been left unused by mistake.

Example

```
int16_t unusedtype ( void )
{
  typedef int16_t local_Type; /* Non-compliant */

  return 67;
}
```

Rule 2.4 A project should not contain unused tag declarations

Category Advisory

Analysis Decidable, System

Applies to C90, C99, C11

Rationale

If a tag is declared but not used, then it is unclear to a reviewer if the tag is redundant or it has been left unused by mistake.

Example

In the following example, the tag `state` is unused and the declaration could have been written without it.

```
void unusedtag ( void )
{
    enum state { S_init, S_run, S_sleep }; /* Non-compliant */
}
```

In the following example, the tag `record_t` is used only in the *typedef* of `record1_t` which is used in the rest of the translation unit whenever the type is needed. This *typedef* can be written in a compliant manner by omitting the tag as shown in the definition of `record2_t`.

```
typedef struct record_t /* Non-compliant */
{
    uint16_t key;
    uint16_t val;
} record1_t;

typedef struct /* Compliant */
{
    uint16_t key;
    uint16_t val;
} record2_t;
```

Rule 2.5 A project should not contain unused macro definitions

Category Advisory

Analysis Decidable, System

Applies to C90, C99, C11

Amplification

#undef of a macro is considered to be a *use* of a macro.

Rationale

If a macro is defined but not used, then it is unclear to a reviewer if the macro is redundant or it has been left unused by mistake.

Example

```
#define SIZE 4
#define DATA 3 /* Non-compliant - DATA not used */

void use_macro ( void )
{
    use_int16 ( SIZE );
}
```

Rule 2.6 A function should not contain unused label declarations

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

If a label is declared but not used, then it is unclear to a reviewer if the label is redundant or it has been left unused by mistake.

Example

```
void unused_label ( void )
{
    int16_t x = 6;
label1:                /* Non-compliant */
    use_int16 ( x );
}
```

Rule 2.7 A function should not contain unused parameters

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

Most functions will be specified as using each of their parameters. If a function parameter is unused, it is possible that the implementation of the function does not match its specification. This rule highlights such potential mismatches.

Example

```
void withunusedpara ( uint16_t *para1,
                    int16_t unusedpara ) /* Non-compliant - unused */
{
    *para1 = 42U;
}
```

Rule 2.8 A project should not contain unused object definitions

Category	Advisory
Analysis	Decidable, System
Applies to	C90, C99, C11

Amplification

An object is unused if the definition (and any declarations) can be removed, and the program still compiles.

Rationale

If an object is defined but unused, then it is unclear to a reviewer if the object is redundant or it has been left unused by mistake.

See also

Rule 8.6

8.3 Comments

Rule 3.1 The character sequences `/*` and `//` shall not be used within a comment

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

If a comment starting sequence, `/*` or `//`, occurs within a `/*` comment, is it quite likely to be caused by a missing `*/` comment ending sequence.

If a comment starting sequence occurs within a `//` comment, it is probably because a region of code has been commented-out using `//`.

Exception

1. *Uniform resource identifiers*, of the form `{scheme}://{path}`, are permitted within comments.
2. The sequence `//` is permitted within a `//` comment.

Example

Consider the following code fragment:

```
/* some comment, end comment marker accidentally omitted
<<New Page>>
Perform_Critical_Safety_Function( X );
/* this comment is non-compliant */
```

In reviewing the page containing the call to the function, the assumption is that it is executed code. Because of the accidental omission of the end comment marker, the call to the safety critical function will not be executed.

In the following C99 example, the presence of `//` comments changes the meaning of the program:

```
x = y // /*
      + z
      // */
      ;
```

This gives `x = y + z`; but would have been `x = y`; in the absence of the two `//` comment start sequences.

The following example demonstrates the use of a URI in a comment, and is compliant by exception 1.

```
/*
** The MISRA C:2012 example suite can be found at
** https://gitlab.com/MISRA/MISRA-C/MISRA-C-2012
*/
```

See also

Dir 4.4

Rule 3.2 Line-splicing shall not be used in // comments

Category Required

Analysis Decidable, Single Translation Unit

Applies to C99, C11

Amplification

Line-splicing occurs when the \ character is immediately followed by a new-line character. If the source file contains multibyte characters, they are converted to the source character set before any splicing occurs.

Rationale

If the source line containing a // comment ends with a \ character in the source character set, the next line becomes part of the comment. This may result in unintentional removal of code.

Example

In the following non-compliant example, the physical line containing the *if* keyword is logically part of the previous line and is therefore a comment.

```
extern bool_t b;

void f ( void )
{
  uint16_t x = 0; // comment \
  if ( b )
  {
    ++x;          /* This is always executed */
  }
}
```

See also

Dir 4.4

8.4 Character sets and lexical conventions

Rule 4.1 Octal and hexadecimal escape sequences shall be terminated

C90 [Implementation 11], C99 [Implementation J.3.4(7, 8)], C11 [Implementation J.3.4(7, 8)]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

An octal or hexadecimal escape sequence shall be terminated by either:

- The start of another escape sequence, or
- The end of the character constant or the end of a string literal.

Rationale

There is potential for confusion if an octal or hexadecimal escape sequence is followed by other characters. For example, the character constant `'\x1f'` consists of a single character whereas the character constant `'\x1g'` consists of the two characters `'\x1'` and `'g'`. The manner in which multi-character constants are represented as integers is implementation-defined.

The potential for confusion is reduced if every octal or hexadecimal escape sequence in a character constant or string literal is terminated.

Example

In this example, each of the strings pointed to by `s1`, `s2` and `s3` is equivalent to "Ag":

```
const char *s1 = "\x41g"; /* Non-compliant */
const char *s2 = "\x41" "g"; /* Compliant - terminated by end of literal */
const char *s3 = "\x41\x67"; /* Compliant - terminated by another escape */

int c1 = '\141t'; /* Non-compliant */
int c2 = '\141\t'; /* Compliant - terminated by another escape */
```

Rule 4.2 Trigraphs should not be used

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

Trigraphs are denoted by a sequence of two question marks followed by a specified third character (e.g. `??~` represents a `~` (tilde) character and `??)` represents a `)`). They can cause accidental confusion with other uses of two question marks.

Note: the so-called digraphs `<:, :>`, `<%, %>`, `%:` and `%;%:` are permitted because they are tokens. Trigraphs are replaced wherever they appear in the program prior to preprocessing.

Example

For example the string:

```
"(Date should be in the form ??-??-??)"
```

would not behave as expected, actually being interpreted by the compiler as:

```
"(Date should be in the form ~~)"
```

8.5 Identifiers

Rule 5.1 *External identifiers* shall be distinct

C90 [Undefined 7], C99 [Unspecified 7; Undefined 28], C11 [Unspecified 8; Undefined 31]

Category	Required
Analysis	Decidable, System
Applies to	C90, C99, C11

Amplification

This rule requires that different *external identifiers* be distinct within the limits imposed by the implementation.

The definition of distinct depends on the implementation and on the edition of the C Standard that is being used:

- In C90 the **minimum** requirement is that the first 6 characters of *external identifiers* are significant but their case is not required to be significant;
- In C99 and later the **minimum** requirement is that the first 31 characters of *external identifiers* are significant, with each universal character or corresponding extended source character occupying between 6 and 10 characters.

In practice, many implementations provide greater limits. For example it is common for *external identifiers* in C90 to be case-sensitive and for at least the first 31 characters to be significant.

Rationale

If two identifiers differ only in non-significant characters, the behaviour is undefined.

If portability is a concern, it would be prudent to apply this rule using the minimum limits specified in the C Standard.

Long identifiers may impair the readability of code. While many automatic code generation systems produce long identifiers, there is a good argument for keeping identifier lengths well below this limit.

Note: In C99 and later, if an extended source character appears in an *external identifier* and that character does not have a corresponding universal character, the C Standard does not specify how many characters it occupies.

Example

In the following example, the definitions all occur in the same translation unit. The implementation in question supports 31 significant case-sensitive characters in *external identifiers*.

```
/*      1234567890123456789012345678901***** Characters */
int32_t engine_exhaust_gas_temperature_raw;
int32_t engine_exhaust_gas_temperature_scaled; /* Non-compliant */

/*      1234567890123456789012345678901***** Characters */
int32_t engine_exhaust_gas_temp_raw;
int32_t engine_exhaust_gas_temp_scaled; /* Compliant */
```

In the following non-compliant example, the implementation supports 6 significant case-insensitive characters in *external identifiers*. The identifiers in the two translation units are different but are not distinct in their significant characters.

```
/* file1.c */
int32_t abc = 0;

/* file2.c */
int32_t ABC = 0;
```

See also

Dir 1.1, Rule 5.2, Rule 5.4, Rule 5.5

Rule 5.2 Identifiers declared in the same *scope* and name space shall be distinct

C90 [Undefined 7], C99 [Undefined 28], C11 [Undefined 31]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule does not apply if:

- both identifiers are *external identifiers*, because this case is covered by Rule 5.1, or
- either identifier is a *macro identifier*, because this case is covered by Rule 5.4 and Rule 5.5.

The definition of distinct depends on the implementation and on the edition of the C Standard that is being used:

- In C90 the **minimum** requirement is that the first 31 characters are significant;
- In C99 and later the **minimum** requirement is that the first 63 characters are significant, with each universal character or extended source character counting as a single character.

Rationale

If two identifiers differ only in non-significant characters, the behaviour is undefined.

If portability is a concern, it would be prudent to apply this rule using the minimum limits specified in the C Standard.

Long identifiers may impair the readability of code. While many automatic code generation systems produce long identifiers, there is a good argument for keeping identifier lengths well below this limit.

Example

In the following example, the implementation in question supports 31 significant case-sensitive characters in identifiers that do not have external linkage.

The identifier `engine_exhaust_gas_temperature_local` is compliant with this rule. Although it is not distinct from the identifier `engine_exhaust_gas_temperature_raw`, it is in a different scope. However, it is not compliant with Rule 5.3.

```
/*          1234567890123456789012345678901*****          Characters */
extern int32_t engine_exhaust_gas_temperature_raw;
static int32_t engine_exhaust_gas_temperature_scaled; /* Non-compliant */

void f ( void )
{
  /*          1234567890123456789012345678901*****          Characters */
  int32_t engine_exhaust_gas_temperature_local; /* Compliant */
}

/*          1234567890123456789012345678901*****          Characters */
static int32_t engine_exhaust_gas_temp_raw;
static int32_t engine_exhaust_gas_temp_scaled; /* Compliant */
```

See also

Dir 1.1, Rule 5.1, Rule 5.3, Rule 5.4, Rule 5.5

Rule 5.3 An identifier declared in an inner scope shall not hide an identifier declared in an outer scope

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

An identifier declared in an inner scope shall be distinct from any identifier declared in an outer scope.

The definition of distinct depends on the implementation and the edition of the C Standard that is being used:

- In C90 the **minimum** requirement is that the first 31 characters are significant;
- In C99 and later the **minimum** requirement is that the first 63 characters are significant, with each universal character or extended source character counting as a single character.

Rationale

If an identifier is declared in an inner scope but is not distinct from an identifier that already exists in an outer scope, then the inner-most declaration will “hide” the outer one. This may lead to developer confusion.

Note: An identifier declared in one name space does not hide an identifier declared in a different name space.

The terms outer and inner scope are defined as follows:

- Identifiers that have file scope can be considered as having the outermost scope;
- Identifiers that have block scope have a more inner scope;
- Successive, nested blocks, introduce more inner scopes.

Example

```
void fn1 ( void )
{
    int16_t i;          /* Declare an object "i"                */

    {
        int16_t i;     /* Non-compliant - hides previous "i"                */
        i = 3;        /* Could be confusing as to which "i" this refers    */
    }
}

struct astruct
{
    int16_t m;
};

extern void g ( struct astruct *p );

int16_t xyz = 0;      /* Declare an object "xyz"                */

void fn2 ( struct astruct xyz ) /* Non-compliant - outer "xyz" is
                               * now hidden by parameter name */
{
    g ( &xyz );
}

uint16_t speed;

void fn3 ( void )
{
    typedef float32_t speed; /* Non-compliant - type hides object */
}
```

See also

Rule 5.2, Rule 5.8

Rule 5.4 *Macro identifiers shall be distinct*

C90 [Undefined 7], C99 [Unspecified 7; Undefined 28], C11 [Unspecified 8; Undefined 31]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule requires that, when a macro is being defined, its name be distinct from:

- The names of the other macros that are currently defined; and
- The names of their parameters.

It also requires that the names of the parameters of a given macro be distinct from each other but does not require that macro parameters names be distinct across two different macros.

The definition of distinct depends on the implementation and on the edition of the C Standard that is being used:

- In C90 the **minimum** requirement is that the first 31 characters of *macro identifiers* are significant;
- In C99 and later the **minimum** requirement is that the first 63 characters of *macro identifiers* are significant.

In practice, implementations may provide greater limits. This rule requires that *macro identifiers* be distinct within the limits imposed by the implementation.

Rationale

If two *macro identifiers* differ only in non-significant characters, the behaviour is undefined. Since macro parameters are active only during the expansion of their macro, there is no issue with parameters in one macro being confused with parameters in another macro.

If portability is a concern, it would be prudent to apply this rule using the minimum limits specified in the C Standard.

Long *macro identifiers* may impair the readability of code. While many automatic code generation systems produce long *macro identifiers*, there is a good argument for keeping *macro identifier* lengths well below this limit.

Note: In C99 and later, if an extended source character appears in a macro name and that character does not have a corresponding universal character, the C Standard does not specify how many characters it occupies.

Example

In the following example, the implementation in question supports 31 significant case-sensitive characters in *macro identifiers*.

```
/*      1234567890123456789012345678901*****          Characters */
#define engine_exhaust_gas_temperature_raw    egt_r
#define engine_exhaust_gas_temperature_scaled  egt_s /* Non-compliant */
```

```

/*      1234567890123456789012345678901*****          Characters */
#define engine_exhaust_gas_temp_raw          egt_r
#define engine_exhaust_gas_temp_scaled      egt_s /* Compliant */

```

See also

Rule 5.1, Rule 5.2, Rule 5.5

Rule 5.5 Identifiers shall be distinct from macro names

C90 [Undefined 7], C99 [Unspecified 7; Undefined 28], C11 [Unspecified 8; Undefined 31]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

This rule requires that the names of macros that exist prior to preprocessing be distinct from the identifiers that exist after preprocessing. It applies to identifiers, regardless of scope or name space, and to any macros that have been defined regardless of whether the definition is still in force when the identifier is declared.

The definition of distinct depends on the implementation and the edition of the C Standard that is being used:

- In C90 the **minimum** requirement is that the first 31 characters are significant;
- In C99 and later the **minimum** requirement is that the first 63 characters are significant, with each universal character or extended source character counting as a single character.

Rationale

Keeping macro names and identifiers distinct can help to avoid developer confusion.

Example

In the following non-compliant example, the name of the *function-like macro* `sum` is also used as an identifier. The declaration of the object `sum` is not subject to macro-expansion because it is not followed by a `(` character. The identifier therefore exists after preprocessing has been performed.

```

#define Sum(x, y) ( ( x ) + ( y ) )

int16_t Sum;

```

The following example is compliant because there is no instance of the identifier `sum` after preprocessing.

```

#define Sum(x, y) ( ( x ) + ( y ) )

int16_t x = Sum ( 1, 2 );

```

In the following example, the implementation in question supports 31 significant case-sensitive characters in identifiers that do not have external linkage. The example is non-compliant because the macro name is not distinct from an identifier name with internal linkage in the first 31 characters.

```
/*          1234567890123456789012345678901***** Characters */
#define     low_pressure_turbine_temperature_1  lp_tb_temp_1
static int32_t low_pressure_turbine_temperature_2;
```

See also

Rule 5.1, Rule 5.2, Rule 5.4

Rule 5.6 A *typedef* name shall be a unique identifier

Category Required

Analysis Decidable, System

Applies to C90, C99, C11

Amplification

A *typedef* name shall be unique across all name spaces and translation units. Multiple declarations of the same *typedef* name are only permitted by this rule if the type definition is made in a *header file* and that *header file* is included in multiple source files.

Rationale

Reusing a *typedef* name either as another *typedef* name or as the name of a function, object or enumeration constant, may lead to developer confusion.

Exception

The *typedef* name may be the same as the structure, union or enumeration tag name associated with the *typedef*.

Example

```
void func ( void )
{
    {
        typedef unsigned char u8_t;
    }
    {
        typedef unsigned char u8_t;    /* Non-compliant - reuse */
    }
}

typedef float mass;

void func1 ( void )
{
    float32_t mass = 0.0f;            /* Non-compliant - reuse */
}

typedef struct list
{
    struct list *next;
    uint16_t     element;
} list;                               /* Compliant - exception */
```

```
typedef struct
{
  struct chain
  {
    struct chain *list;
    uint16_t     element;
  } s1;
  uint16_t length;
} chain;                                     /* Non-compliant - tag "chain" not
                                           * associated with typedef */
```

See also

Rule 5.7

Rule 5.7 A tag name shall be a unique identifier

Category Required

Analysis Decidable, System

Applies to C90, C99, C11

Amplification

The tag shall be unique across all name spaces and translation units.

All declarations of the tag shall specify the same type.

Multiple complete declarations of the same tag are only permitted by this rule if the tag is declared in a *header file* and that *header file* is included in multiple source files.

Rationale

Reusing a tag name may lead to developer confusion.

There is also undefined behaviour associated with reuse of tag names in C90, although this is not listed in that C Standard's Annex. This is a *constraint* violation in C99 and later.

Exception

The tag name may be the same as the *typedef* name with which it is associated.

Example

```
struct stag
{
  uint16_t a;
  uint16_t b;
};

struct stag a1 = { 0, 0 }; /* Compliant - compatible with above */
union stag a2 = { 0, 0 }; /* Non-compliant - declares different type
                          * from struct stag
                          * Constraint violation in C99 */
```

The following example also violates Rule 5.3

```

struct deer
{
    uint16_t a;
    uint16_t b;
};

void foo ( void )
{
    struct deer
    {
        uint16_t a;
    };
    /* Non-compliant - tag "deer" reused */
}

typedef struct coord
{
    uint16_t x;
    uint16_t y;
} coord;
/* Compliant by Exception */

struct elk
{
    uint16_t x;
};

struct elk /* Non-compliant - declaration of different type
           * Constraint violation in C99 */
{
    uint32_t x;
};

```

See also

Rule 5.6

Rule 5.8 Identifiers that define objects or functions with external linkage shall be unique

Category	Required
Analysis	Decidable, System
Applies to	C90, C99, C11

Amplification

An identifier used as an *external identifier* shall not be used for any other purpose in any name space or translation unit, even if it denotes an object with no linkage.

Rationale

Enforcing uniqueness of identifier names in this manner helps avoid confusion. Identifiers of objects that have no linkage need not be unique since there is minimal risk of such confusion.

Example

In the following example, `file1.c` and `file2.c` are both part of the same project.

```
/* file1.c */
int32_t count;          /* "count" has external linkage */
void foo ( void )      /* "foo" has external linkage */
{
    int16_t index;      /* "index" has no linkage */
}

/* file2.c */
static void foo ( void ) /* Non-compliant - "foo" is not unique
* (it is already defined with external
* linkage in file1.c) */
{
    int16_t count;      /* Non-compliant - "count" has no linkage
* but clashes with an identifier with
* external linkage */
    int32_t index;      /* Compliant - "index" has no linkage */
}
```

See also

Rule 5.3

Rule 5.9 Identifiers that define objects or functions with internal linkage should be unique

Category Advisory

Analysis Decidable, System

Applies to C90, C99, C11

Amplification

An identifier name that defines objects or functions with internal linkage should be unique across all name spaces and translation units. Any identifier used in this way should not have the same name as any other identifier, even if that other identifier denotes an object with no linkage.

Rationale

Enforcing uniqueness of identifier names in this manner helps avoid confusion.

Exception

An *inline function* with internal linkage may be defined in more than one translation unit provided that all such definitions are made in the same *header file* that is included in each translation unit.

Example

In the following example, `file1.c` and `file2.c` are both part of the same project.

```
/* file1.c */
static int32_t count;    /* "count" has internal linkage */
static void foo ( void ) /* "foo" has internal linkage */
{
    int16_t count;      /* Non-compliant - "count" has no linkage
* but clashes with an identifier with
* internal linkage */
    int16_t index;      /* "index" has no linkage */
}
```

```

void bar1 ( void )
{
    static int16_t count;    /* Non-compliant - "count" has no linkage
                           * but clashes with an identifier with
                           * internal linkage                               */
    int16_t index;         /* Compliant - "index" is not unique but
                           * has no linkage                               */
    foo ( );
}

/* End of file1.c */

/* file2.c */
static int8_t count;      /* Non-compliant - "count" has internal
                           * linkage but clashes with other
                           * identifiers of the same name                */
static void foo ( void ) /* Non-compliant - "foo" has internal
                           * linkage but clashes with a function of
                           * the same name                                */
{
    int32_t index;        /* Compliant - both "index" and "nbytes"
                           * are not unique but have no linkage          */
    int16_t nbytes;
}

void bar2 ( void )
{
    static uint8_t nbytes; /* Compliant - "nbytes" is not unique but
                           * has no linkage and the storage class is
                           * irrelevant                                   */
}

/* End of file2.c */

```

See also

Rule 8.10

8.6 Types

Rule 6.1 Bit-fields shall only be declared with an appropriate type

C90 [Undefined 38; Implementation 29]
 C99 [Implementation J.3.9(1, 2)]
 C11 [Implementation J.3.9(1, 2)]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

The appropriate bit-field types are:

- C90: either *unsigned int* or *signed int*;
- C99 and later: one of:
 - either *unsigned int* or *signed int*;
 - another explicitly signed or explicitly unsigned integer type that is permitted by the implementation;
 - `_Bool`.

Note: It is permitted to use *typedefs* to designate an appropriate type.

Rationale

Using *int* is implementation-defined because bit-fields of type *int* can be either *signed* or *unsigned*.

The use of *enum*, *short*, *char* or any other type for bit-fields is not permitted in C90 because the behaviour is undefined.

In C99 and later, the implementation may define other integer types that are permitted in bit-field declarations.

Example

The following example is applicable to implementations that do not provide any additional bit-field types. It assumes that the *int* type is 16-bit.

```
typedef unsigned int UINT_16;

struct s {
    unsigned int b1:2; /* Compliant */
    int          b2:2; /* Non-compliant - plain int not permitted */
    UINT_16      b3:2; /* Compliant - typedef designating unsigned int */
    signed long  b4:2; /* Non-compliant even if long and int are the
                       * same size */
};
```

Rule 6.2 Single-bit named bit-fields shall not be of a signed type

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

For C99 and later, the C Standard states that signed integers have exactly one sign-bit, meaning that a single-bit signed bit-field will have no value-bits. In any representation of integers, a meaningful value cannot be specified if there are zero value-bits.

A single-bit signed bit-field is therefore unlikely to behave in a useful way and its presence is likely to indicate programmer confusion.

Although C90 does not provide so much detail regarding the representation of types, the same considerations apply.

Note: this rule does not apply to unnamed bit-fields as their values cannot be accessed.

Rule 6.3 A bit field shall not be declared as a member of a union

C90 [Implementation 30]
 C99 [Implementation J.3.9(4)]
 C11 [Implementation J.3.9(5)]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

A member of a union shall not be declared as a bit field.

This rule does not apply to sub-objects within union members that do not themselves have a union type.

Rationale

The exact bitwise position of a bit field within a storage unit is implementation defined. Therefore, if two bit fields are declared such that they fit within the same storage unit of a union, the compiler is not required to overlay them over one another beginning from the starting bit of the storage unit.

If the union is used for *type-punning*, it is therefore unclear which bits of the previously-stored value will be accessed by the bit field.

If the union is not intended to be used for *type-punning*, there is no point in declaring the members as bit fields, because no space will be saved (a complete storage unit will need to be allocated within the union anyway).

Example

```

/* Compliant      - if the user wants to type-pun, the bits of 'big' which will
                   ... be overlaid are clearly identified */
union U1 {
    uint8_t small;
    uint32_t big;
};

/* Non-compliant - it is unclear which, if any, bits of 'big' are overlaid by
                   ... 'small' in this type */
union U2 {
    uint32_t small:8;
    uint32_t big;
};

/* Non-compliant - it is unclear if any bits of 'big' are overlaid by
                   ... 'small' in this type */
union U3 {
    uint32_t small: 8;
    uint32_t big  :24;
};

/* Compliant      - a sub-object can be a bit-field */
union U4 {
    struct {
        uint8_t a:4;
        uint8_t b:4;
        uint8_t c:4;
        uint8_t d:4;
    } q;
    uint16_t r;
};

```

8.7 Literals and constants

Rule 7.1 Octal constants shall not be used

[Koenig 9]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

Developers writing constants that have a leading zero might expect them to be interpreted as decimal constants.

Note: this rule does not apply to octal escape sequences because the use of a leading `\` character means that there is less scope for confusion.

Exception

The integer constant zero (written as a single numeric digit), is strictly speaking an octal constant, but is a permitted exception to this rule.

Example

```
extern uint16_t code[ 10 ];

code[ 1 ] = 109; /* Compliant - decimal 109 */
code[ 2 ] = 100; /* Compliant - decimal 100 */
code[ 3 ] = 052; /* Non-Compliant - decimal 42 */
code[ 4 ] = 071; /* Non-Compliant - decimal 57 */
```

Rule 7.2 A “u” or “U” suffix shall be applied to all integer constants that are represented in an unsigned type

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

This rule applies to:

- Integer constants that appear in the controlling expressions of `#if` and `#elif` preprocessing directives;
- Any other integer constants that exist after preprocessing.

Note: during preprocessing, the type of an integer constant is determined in the same manner as after preprocessing except that:

- All signed integer types behave as if they were *long* (C90) or *intmax_t* (C99 and later);
- All unsigned integer types behave as if they were *unsigned long* (C90) or *uintmax_t* (C99 and later).

Rationale

The type of an integer constant is a potential source of confusion, because it is dependent on a complex combination of factors including:

- The magnitude of the constant;
- The implemented sizes of the integer types;
- The presence of any suffixes;
- The number base in which the value is expressed (i.e. decimal, octal or hexadecimal).

For example, the integer constant 40000 is of type *signed int* in a 32-bit environment but of type *signed long* in a 16-bit environment. The value 0x8000 is of type *unsigned int* in a 16-bit environment, but of type *signed int* in a 32-bit environment.

Note:

- Any value with a “U” suffix is of unsigned type;
- An unsuffixed decimal value less than 2^{31} is of signed type.

But:

- An unsuffixed hexadecimal value greater than or equal to 2^{15} may be of signed or unsigned type;
- For C90, an unsuffixed decimal value greater than or equal to 2^{31} may be of signed or unsigned type.

Signedness of constants should be explicit. If a constant is of an unsigned type, applying a “U” suffix makes it clear that the programmer understands that the constant is unsigned.

Note: this rule does not depend on the context in which a constant is used; promotion and other conversions that may be applied to the constant are not relevant in determining compliance with this rule.

Example

The following example assumes a machine with a 16-bit *int* type and a 32-bit *long* type. It shows the type of each integer constant determined in accordance with the C Standard. The integer constant `0x8000` is non-compliant because it has an unsigned type but does not have a “U” suffix.

Constant	Type	Compliance
<code>32767</code>	<i>signed int</i>	Compliant
<code>0x7fff</code>	<i>signed int</i>	Compliant
<code>32768</code>	<i>signed long</i>	Compliant
<code>32768u</code>	<i>unsigned int</i>	Compliant
<code>0x8000</code>	<i>unsigned int</i>	Non-compliant
<code>0x8000u</code>	<i>unsigned int</i>	Compliant

Rule 7.3 The lowercase character “l” shall not be used in a literal suffix

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

Using the uppercase suffix “L” removes the potential ambiguity between “1” (digit 1) and “l” (letter “el”) when declaring literals.

Example

Note: the examples containing the *long long* suffix are applicable only to C99.

```
const int64_t a = 0L;
const int64_t b = 0l; /* Non-compliant */
const uint64_t c = 0Lu;
const uint64_t d = 0lU; /* Non-compliant */
const uint64_t e = 0ULL;
const uint64_t f = 0Ull; /* Non-compliant */
const int128_t g = 0LL;
const int128_t h = 0ll; /* Non-compliant */
const float128_t m = 1.2L;
const float128_t n = 2.4l; /* Non-compliant */
```

Rule 7.4 A string literal shall not be *assigned* to an object unless the object's type is "pointer to *const*-qualified *char*"

C90 [Undefined 12], C99 [Unspecified 14; Undefined 30], C11 [Unspecified 15; Undefined 33]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The type used to represent characters within a string literal depends on the encoding prefix:

- *char* when there is no prefix — *character string literal*;
- *char* for a `u8` prefix — *UTF-8 string literal*;
- *wchar_t* for an `L` prefix — *wide string literal*;
- *char16_t* for a `u` prefix — *wide string literal*;
- *char32_t* for a `U` prefix — *wide string literal*.

No attempt shall be made to directly modify a character string literal, a UTF-8 string literal or a wide string literal.

The result of the address-of operator, `&`, applied to a character string literal or a UTF-8 string literal shall not be *assigned* to an object unless that object's type is "pointer to array of *const*-qualified *char*".

The result of the address-of operator, `&`, applied to a wide string literal shall not be *assigned* to an object unless that object's type is "pointer to array of *const*-qualified *wchar_t*", "pointer to array of *const*-qualified *char16_t*" or "pointer to array of *const*-qualified *char32_t*", as appropriate for the encoding prefix.

Rationale

Any attempt to modify a string literal results in undefined behaviour. For example, some implementations may store string literals in read-only memory in which case an attempt to modify the string literal will fail and may also result in an exception or crash.

This rule, when applied in conjunction with others, prevents a string literal from being modified.

It is explicitly unspecified in C99 and later whether string literals that share a common ending are stored in distinct memory locations. Therefore, even if an attempt to modify a string literal appears to succeed, it is possible that another string literal might be inadvertently altered.

Exception

This rule does not apply to a string literal passed as an argument to the variable argument list of a variadic function.

Example

The following example shows an attempt to modify a string literal directly:

```
"0123456789"[0] = '*'; /* Non-compliant */
```

These examples show how to prevent modification of string literals indirectly:

```
/* Non-compliant - s is not const-qualified */
char *s = "string";

/* Compliant - p is const-qualified; additional qualifiers are permitted */
const volatile char *p = "string";

extern void f1 ( char *s1 );

extern void f2 ( const char *s2 );

void g ( void )
{
    f1 ( "string" ); /* Non-compliant - parameter s1 is not
                    * const-qualified */
    f2 ( "string" ); /* Compliant */
}

char *name1 ( void )
{
    return ( "MISRA" ); /* Non-compliant - return type is not
                        * const-qualified */
}

const char *name2 ( void )
{
    return ( "MISRA" ); /* Compliant */
}
```

This example shows the permitted exemption for variadic functions:

```
extern void f3( uint16_t x, ... ); /* Note: non-compliant with Rule 17.1 */
extern void f4( char *text, ... ); /* Note: non-compliant with Rule 17.1 */

void variadic( void )
{
    f3( 42u, "MISRA" ); /* Compliant by exception */
    f4( "MISRA", 42u ); /* Non-compliant - exception only applies to
                        variable argument lists */
}
```

See also

Rule 11.4, Rule 11.8, Rule 17.1

Rule 7.5 The argument of an integer constant macro shall have an appropriate form

C99 [Undefined 137], C11 [Undefined 145]

Category	Mandatory
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Amplification

The argument of an integer constant macro shall satisfy the following:

- The argument must be an unsuffixed integer (decimal, octal or hexadecimal) literal;
- The value of the argument must not exceed the limits for the equivalent *exact-width* type indicated by the name of the macro used. For example, the argument to `UINT16_C` must be representable as an unsigned 16-bit value.

Rationale

The behaviour is undefined if the argument of an integer constant macro does not have an appropriate form.

Example

```
#include <stdint.h>

uint32_t u1 = UINT32_C( 10 ); /* Compliant */
uint32_t u2 = UINT32_C( 10UL ); /* Non-compliant - constant is suffixed */
uint32_t u3 = UINT32_C( 10.0 ); /* Non-compliant - floating-point constant */
uint16_t u4 = UINT16_C( -2 ); /* Non-compliant - constant expression */

int32_t s1 = INT32_C( -2 ); /* Non-compliant - constant expression */
int32_t s2 = -INT32_C( 2 ); /* Compliant */
```

In the following example, the constant has a value that cannot be represented in 16 bits.

```
uint_least16_t u5 = UINT16_C ( 0x10000 ); /* Non-compliant, even if uint_least16_t
... is implemented as a 32-bit type */
```

See also

Rule 7.6

Rule 7.6 The small integer variants of the minimum-width integer constant macros shall not be used

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Amplification

The minimum-width integer constant macros are of the form `INTn_C(value)` and `UINTn_C(value)`, where *n* is a value corresponding to a type `int_leastn_t`.

Small integer refers to any integer type with width less than that of type `int`.

Rationale

The C Standard requires that the minimum-width integer constant macros expand to an integer constant expression suitable for use in `#if` pre-processing directive, and that the type of the expression has the same type as would result from integer promotion. Consequentially many implementations of the small integer macros have opted to simply substitute the macro for the argument. This results in an expression with type `int` and not the type that may have been anticipated by the use of the macro.

Example

```
int main( void )
{
    uint8_t a = UINT8_C( 100 ); /* Non-compliant - typically expands as plain 100
                                i.e. as a signed int                */
}
```

The following example shows the impact of the typing conflict:

```
#define M(x) _Generic( (x), uint8_t: fu8, default: fi )(x)

int main( void )
{
    M( UINT8_C( 100 ) ); /* Non-compliant - selects fi, not fu8 */
}
```

See also

Rule 7.5

8.8 Declarations and definitions

Rule 8.1 Types shall be explicitly specified

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90

Rationale

The C90 standard permits types to be omitted in some circumstances, in which case the *int* type is implicitly specified. Examples of the circumstances in which an implicit *int* might be used are:

- Object declarations;
- Parameter declarations;
- Member declarations;
- *typedef* declarations;
- Function return types.

The omission of an explicit type might lead to confusion. For example, in the declaration:

```
extern void g ( char c, const k );
```

the type of *k* is *const int* whereas *const char* might have been expected.

Example

The following examples show compliant and non-compliant object declarations:

```
extern          x;      /* Non-compliant - implicit int type */
extern int16_t x;      /* Compliant - explicit type          */
const          y;      /* Non-compliant - implicit int type */
const int16_t y;      /* Compliant - explicit type          */
```

The following examples show compliant and non-compliant function type declarations:

```
extern f ( void );                /* Non-compliant - implicit
                                * int return type          */
extern int16_t f ( void );        /* Compliant                */

extern void g ( char c, const k ); /* Non-compliant - implicit
                                * int for parameter k      */
extern void g ( char c, const int16_t k ); /* Compliant                */
```

The following examples show compliant and non-compliant type definitions:

```
typedef ( *pfi ) ( void );        /* Non-compliant - implicit int
                                * return type                */
typedef int16_t ( *pfi ) ( void ); /* Compliant                */
typedef void ( *pfv ) ( const x ); /* Non-compliant - implicit int
                                * for parameter x                */
typedef void ( *pfv ) ( int16_t x ); /* Compliant                */
```

The following examples show compliant and non-compliant member declarations:

```
struct str
{
    int16_t x;    /* Compliant */
    const y;    /* Non-compliant - implicit int for member y */
} s;
```

See also

Rule 8.2

Rule 8.2 Function types shall be in *prototype form* with named parameters

C90 [Undefined 22–25], C99 [Undefined 36–39, 73, 79], C11 [Undefined 38–41, 79, 85]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

The early edition of C, commonly referred to as K&R C [47], did not provide a mechanism for checking the number of arguments or their types against the corresponding parameters. The type of an object or function did not have to be declared in K&R C since the default type of an object and the default return type of a function was *int*.

The C90 Standard introduced function prototypes, a form of function declarator in which the parameter types were declared. This permitted argument types to be checked against parameter types. It also allowed the number of arguments to be checked except when a function prototype specified that a variable number of arguments was expected. The C90 standard did not **require** the use of function prototypes for reasons of backward compatibility with existing code. For the same reason, it continued to permit types to be omitted in which case the type would default to *int*.

The C99 Standard removed the default *int* type from the language but continued to allow K&R-style function types in which there was no means to supply parameter type information in a declaration and it was optional to supply parameter type information in a definition.

The mismatch between the number of arguments and parameters, their types and the expected and actual return type of a function provides potential for undefined behaviour. The purpose of this rule along with Rule 8.1 and Rule 8.4 is to avoid this undefined behaviour by requiring parameter types and function return types to be specified explicitly. Rule 17.3 ensures that this information is available at the time of a function call, thereby requiring the compiler to diagnose any mismatch that is detected.

This rule also requires that names be specified for all the parameters in a declaration. The parameter names can provide useful information regarding the function interface and a mismatch between a declaration and definition might be indicative of a programming error.

Note: An empty parameter list is **not** valid in a prototype. If a function type has no parameters its *prototype form* uses the keyword *void*.

Example

The first example shows declarations of some functions and the corresponding definitions for some of those functions.

```

/* Compliant */
extern int16_t func1 ( int16_t n );

/* Non-compliant - parameter name not specified */
extern void func2 ( int16_t );

/* Non-compliant - not in prototype form */
static int16_t func3 ( );

/* Compliant - prototype specifies 0 parameters */
static int16_t func4 ( void );

/* Compliant */
int16_t func1 ( int16_t n )
{
    return n;
}

/* Non-compliant - old style identifier and declaration list */
static int16_t func3 ( vec, n )
int16_t *vec;
int16_t n;
{
    return vec[ n - 1 ];
}

```

This example section shows the application of the rule to function types other than in function declarations and definitions.

```

/* Non-compliant - no prototype */
int16_t ( *pf1 ) ( );

/* Compliant - prototype specifies 0 parameters */
int16_t ( *pf1 ) ( void );

/* Non-compliant - parameter name not specified */
typedef int16_t ( *pf2_t ) ( int16_t );

/* Compliant */
typedef int16_t ( *pf3_t ) ( int16_t n );

```

See also

Rule 1.5, Rule 8.1, Rule 8.3, Rule 8.4, Rule 17.3

Rule 8.3 All declarations of an object or function shall use the same names and type qualifiers

C90 [Undefined 10], C99 [Undefined 14], C11 [Undefined 15], [Koenig 59–62]

Category Required

Analysis Decidable, System

Applies to C90, C99, C11

Amplification

Storage class specifiers are not included within the scope of this rule.

Rationale

Using types and qualifiers consistently across declarations of the same object or function encourages stronger typing.

Specifying parameter names in function prototypes allows the function definition to be checked for interface consistency with its declarations.

Exception

1. Compatible versions of the same basic type may be used interchangeably. For example, *int*, *signed* and *signed int* are all equivalent.
2. The naming requirements of this rule do not apply to unnamed function parameters. This is covered by Rule 8.2.

Example

```
extern void f ( signed int a );
void f ( int a ); /* Compliant - Exception 1 */

extern void g ( signed int b );
extern void g ( signed int ); /* Compliant - Exception 2 */

extern void h ( int * const c );
extern void h ( int * c ); /* Non-compliant - mis-matched type qualifiers */

extern void j ( int d );
extern void j ( int e ); /* Non-compliant - mis-matched parameter names */
```

Note: all the above are not compliant with Dir 4.6; example `g()` is also not compliant with Rule 8.2.

```
extern int16_t func ( int16_t num, int16_t den );

/* Non-compliant - parameter names do not match */
int16_t func ( int16_t den, int16_t num )
{
    return num / den;
}
```

In this example the definition of `area` uses a different type name for the parameter `h` from that used in the declaration. This does not comply with the rule even though `width_t` and `height_t` are the same basic type.

```
typedef uint16_t width_t;
typedef uint16_t height_t;
typedef uint32_t area_t;

extern area_t area ( width_t w, height_t h );

area_t area ( width_t w, width_t h )
{
    return ( area_t ) w * h;
}
```

This rule does not require that a function pointer declaration use the same names as a function declaration. The following example is therefore compliant.

```
extern void f1 ( int16_t x );
extern void f2 ( int16_t y );

void f ( bool_t b )
{
    void ( *fp1 ) ( int16_t z ) = b ? f1 : f2;
}
```

See also

Rule 8.2, Rule 8.4

Rule 8.4 A compatible declaration shall be visible when an object or function with external linkage is defined

C90 [Undefined 24], C99 [Undefined 39], C11 [Undefined 41]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

A compatible declaration is one which declares a compatible type for the object or function being defined.

Rationale

If a declaration for an object or function is visible when that object or function is defined, a compiler must check that the declaration and definition are compatible. In the presence of function prototypes, as required by Rule 8.2, checking extends to the number and type of function parameters.

The recommended method of implementing declarations of objects and functions with external linkage is to declare them in a *header file*, and then include the *header file* in all those code files that need them, including the one that defines them (See Rule 8.5).

Exception

The function `main` need not have a separate declaration.

Example

In these examples there are no declarations or definitions of objects or functions other than those present in the code.

```
extern int16_t count;
    int16_t count = 0;          /* Compliant          */

extern uint16_t speed = 6000u; /* Non-compliant - no declaration
    * prior to this definition */

uint8_t pressure = 101u;     /* Non-compliant - no declaration
    * prior to this definition */
```

```
extern void func1 ( void );
extern void func2 ( int16_t x, int16_t y );
extern void func3 ( int16_t x, int16_t y );

void func1 ( void )
{
    /* Compliant */
}
```

The following non-compliant definition of `func3` also violates Rule 8.3.

```
void func2 ( int16_t x, int16_t y )
{
    /* Compliant */
}

void func3 ( int16_t x, uint16_t y )
{
    /* Non-compliant - parameter types different */
}

void func4 ( void )
{
    /* Non-compliant - no declaration of func4 before this definition */
}

static void func5 ( void )
{
    /* Compliant - rule does not apply to objects/functions with internal
    * linkage */
}
```

See also

Rule 8.2, Rule 8.3, Rule 8.5, Rule 17.3

Rule 8.5 An external object or function shall be declared once in one and only one file

[Koenig 66]

Category Required

Analysis Decidable, System

Applies to C90, C99, C11

Amplification

This rule applies to non-defining declarations only.

Rationale

Typically, a single declaration will be made in a *header file* that will be included in any translation unit in which the identifier is defined or used. This ensures consistency between:

- The declaration and the definition;
- Declarations in different translation units.

Note: there may be many *header files* in a project, but each external object or function shall only be declared in one *header file*.

Example

```
/* featureX.h */
extern int16_t a;    /* Declare a */

/* file.c */
#include "featureX.h"

int16_t a = 0;      /* Define a */
```

See also

Rule 8.4

Rule 8.6 An identifier with external linkage shall have exactly one external definition

C90 [Undefined 44], C99 [Undefined 78], C11 [Undefined 84], [Koenig 55, 63–65]

Category Required

Analysis Decidable, System

Applies to C90, C99, C11

Rationale

The behaviour is undefined if an identifier is used for which multiple definitions exist (in different files) or no definition exists at all. Multiple definitions in different files are not permitted by this rule even if the definitions are the same. It is undefined behaviour if the declarations are different, or initialize the identifier to different values.

Example

In this example the object `i` is defined twice.

```
/* file1.c */
int16_t i = 10;

/* file2.c */
int16_t i = 20;    /* Non-compliant - two definitions of i */
```

In this example the object `j` has one tentative definition and one external definition.

```
/* file3.c */
int16_t j;        /* Tentative definition */
int16_t j = 1;    /* Compliant - external definition */
```

The following example is non-compliant because the object `k` has two external definitions. The tentative definition in `file4.c` becomes an external definition at the end of the translation unit.

```
/* file4.c */
int16_t k;        /* Tentative definition - becomes external */

/* file5.c */
int16_t k = 0;    /* External definition */
```

See also

Rule 8.15

Rule 8.7 Functions and objects should not be defined with external linkage if they are referenced in only one translation unit

[Koenig 56, 57]

Category	Advisory
Analysis	Decidable, System
Applies to	C90, C99, C11

Rationale

Restricting the visibility of an object by giving it internal linkage or no linkage reduces the chance that it might be accessed inadvertently. Similarly, reducing the visibility of a function by giving it internal linkage reduces the chance of it being called inadvertently.

Compliance with this rule also avoids any possibility of confusion between an identifier and an identical identifier in another translation unit or a library.

Example

```

/* file.h */
extern void ext_fn1 ( void ); /* Compliant */
extern void ext_fn2 ( void ); /* Non-compliant */

/* file1.c */
#include "file.h"
void ext_fn1 ( void ) /* Compliant - defined in this translation unit,
                        but used externally */
{
    /* Function definition */
}

void ext_fn2 ( void ) /* Non-compliant - defined and used only
                        in this translation unit */
{
    /* Function definition */
}

void fn_file1 ( void )
{
    ext_fn2( );
}

/* file2.c */
#include "file.h"
void fn_file2 ( void )
{
    ext_fn1( );
}

```

Rule 8.8 The *static* storage class specifier shall be used in all declarations of objects and functions that have internal linkage

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

Since definitions are also declarations, this rule applies equally to definitions.

Rationale

The C Standard states that if an object or function is declared with the *extern* storage class specifier and another declaration of the object or function is already visible, the linkage is that specified by the earlier declaration. This can be confusing because it might be expected that the *extern* storage class specifier creates external linkage. The *static* storage class specifier shall therefore be consistently applied to objects and functions with internal linkage.

Example

```
static int32_t x = 0;          /* definition: internal linkage */
extern int32_t x;            /* Non-compliant */

static int32_t f ( void );    /* declaration: internal linkage */
int32_t f ( void )           /* Non-compliant */
{
    return 1;
}

static int32_t g ( void );    /* declaration: internal linkage */
extern int32_t g ( void )     /* Non-compliant */
{
    return 1;
}
```

See also

Rule 1.5

Rule 8.9 An object should be declared at block scope if its identifier only appears in a single function

Category	Advisory
Analysis	Decidable, System
Applies to	C90, C99, C11

Rationale

Declaring an object at block scope reduces the possibility that the object might be accessed inadvertently and makes clear the intention that it should not be accessed elsewhere.

Within a function, whether objects are declared at the outermost or innermost block is largely a matter of style.

It is recognized that there are situations in which it may not be possible to comply with this rule. For example, an object with static storage duration declared at block scope cannot be accessed directly from outside the block. This makes it impossible to set up and check the results of unit test cases without using indirect accesses to the object. In this kind of situation, some projects may prefer not to apply this rule.

Example

In this compliant example, `i` is declared at block scope because it is a *loop counter*. There is no need for other functions in the same file to use the same object for any other purpose.

```
void func ( void )
{
    int32_t i;

    for ( i = 0; i < N; ++i )
    {
    }
}
```

In this compliant example, the function `count` keeps track of the number of times it has been called and returns that number. No other function needs to know the details of the implementation of `count` so the call counter is declared with block scope.

```
uint32_t count ( void )
{
    static uint32_t call_count = 0;

    ++call_count;
    return call_count;
}
```

Rule 8.10 An *inline function* shall be declared with the static storage class

C99 [Unspecified 20; Undefined 67], C99 [Unspecified 21; Undefined 70]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Rationale

If an *inline function* is declared with external linkage but not defined in the same translation unit, the behaviour is undefined.

A call to an *inline function* declared with external linkage may call the external definition of the function, or it may use the inline definition. Although this should not affect the behaviour of the called function, it might affect execution timing and therefore have an impact on a real-time program.

Note: an *inline function* can be made available to several translation units by placing its definition in a *header file*.

See also

Rule 5.9

Rule 8.11 When an array with external linkage is declared, its size should be explicitly specified

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule applies to non-defining declarations only. It is possible to define an array and specify its size implicitly by means of initialization.

Rationale

Although it is possible to declare an array with incomplete type and access its elements, it is safer to do so when the size of the array may be explicitly determined. Providing size information for each declaration permits them to be checked for consistency. It may also permit a static checker to perform some array bounds analysis without needing to analyse more than one translation unit.

Example

```
extern int32_t array1[ 10 ];    /* Compliant      */
extern int32_t array2[ ];     /* Non-compliant */
```

Rule 8.12 Within an enumerator list, the value of an implicitly-specified enumeration constant shall be unique

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

An implicitly-specified enumeration constant has a value 1 greater than its predecessor. If the first enumeration constant is implicitly-specified then its value is 0.

An explicitly-specified enumeration constant has the value of the associated constant expression.

If implicitly-specified and explicitly-specified constants are mixed within an enumeration list, it is possible for values to be replicated. Such replication may be unintentional and may give rise to unexpected behaviour.

This rule requires that any replication of enumeration constants be made explicit, thus making the intent clear.

Example

In the following examples the `green` and `yellow` enumeration constants are given the same value.

```
/* Non-compliant - yellow replicates implicit green */
enum colour { red = 3, blue, green, yellow = 5 };

/* Compliant
enum colour { red = 3, blue, green = 5, yellow = 5 };
```

Rule 8.13 A pointer should point to a *const*-qualified type whenever possible

Category	Advisory
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

A pointer should point to a *const*-qualified type unless either:

- It is used to modify an object, or
- It is copied to another pointer that points to a type that is not *const*-qualified by means of either:
 - *Assignment*, or
 - Memory move or copying functions.

For the purposes of simplicity, this rule is written in terms of pointers and the types that they point to. However, it applies equally to arrays and the types of the elements that they contain. An array should have elements with *const*-qualified type unless either:

- Any element of the array is modified, or
- It is copied to a pointer that points to a type that is not *const*-qualified by the means described above.

Rationale

This rule encourages best practice by ensuring that pointers are not inadvertently used to modify objects. Conceptually, it is equivalent to initially declaring:

- All arrays to have elements with *const*-qualified type, and
- All pointers to point to *const*-qualified types.

and then removing *const*-qualification only where it is necessary to comply with the *constraints* of the language standard.

Example

In the following non-compliant example, *p* is not used to modify an object but the type to which it points is not *const*-qualified.

```
uint16_t f ( uint16_t *p )
{
    return *p;
}
```

The code would be compliant if the function were defined with:

```
uint16_t g ( const uint16_t *p )
```

The following example violates a *constraint* because an attempt is made to use a *const*-qualified pointer to modify an object.

```
void h ( const uint16_t *p )
{
    *p = 0;
}
```

In the following example, the pointer *s* is *const*-qualified but the type it points to is not. Since *s* is not used to modify an object, this is non-compliant.

```
#include <string.h>

char last_char ( char * const s )
{
    return s[ strlen ( s ) - 1u ];
}
```

The code would be compliant if the function were defined with:

```
char last_char ( const char * const s )
```

In this non-compliant example, none of the elements of the array *a* are modified but the element type is not *const*-qualified.

```
uint16_t first ( uint16_t a[ 5 ] )
{
    return a[ 0 ];
}
```

The code would be compliant if the function were defined with:

```
uint16_t first ( const uint16_t a[ 5 ] )
```

Rule 8.14 The *restrict* type qualifier shall not be used

C99 [Undefined 65, 66], C11 [Undefined 68, 69]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Rationale

When used with care the *restrict* type qualifier may improve the efficiency of code generated by a compiler. It may also allow improved static analysis. However, to use the *restrict* type qualifier the programmer must be sure that the memory areas operated on by two or more pointers do not overlap.

There is a significant risk that a compiler will generate code that does not behave as expected if *restrict* is used incorrectly.

Example

The following example is compliant because the MISRA C Guidelines do not apply to the Standard Library functions. The programmer must ensure that the areas defined by *p*, *q* and *n* do not overlap.

```
#include <string.h>

void f ( void )
{
    /* memcpy has restrict-qualified parameters */
    memcpy ( p, q, n );
}
```

The following example is non-compliant because a function has been defined using *restrict*.

```
void user_copy ( void * restrict p, void * restrict q, size_t n )
{
}
```

Rule 8.15 All declarations of an object with an explicit *alignment specification* shall specify the same *alignment*

C11 [Undefined 73]

Category	Required
Analysis	Decidable, System
Applies to	C11

Amplification

If any declaration (including the definition) of an object includes an explicit *alignment specification*, all other declarations of the same object shall also include the same explicit *alignment specification*.

Two *alignment specifications* are the same if they are lexically identical after macro expansion.

This rule applies both to alignments specified directly by an expression, and by naming a type.

Rationale

If multiple declarations of an object with external linkage have conflicting *alignment specifications*, the behaviour is undefined. Therefore, if the alignment of the object is important, it should be listed explicitly in order to avoid the risk of creating conflicting declarations in multiple translation units. If the alignment of the object does not need to be set explicitly, the *alignment specification* should be omitted entirely to avoid confusion.

Since an object with internal linkage is only accessible by name from within a single translation unit, there is no direct risk of incompatible alignment specifications causing undefined behaviour. However, this rule reduces the risk of accidentally introducing undefined behaviour during maintenance if the code is later modified to give an object external linkage instead.

If the alignment operand is a type name and the *alignment specifications* do not consistently use that same type name, there is a risk of introducing inconsistency if the configuration changes, such as if code is recompiled on a different platform with different fundamental type alignments.

Example

```

/* header.h - #included by both file1.c and file2.c */

extern alignas (16)          int32_t a;
extern alignas (0)          int32_t b;
extern                      int32_t c;
extern                      int32_t d;
extern alignas (16)          int32_t e;
extern alignas (16)          int32_t f;
extern                      int32_t g;

extern alignas (float)       int32_t i;
extern alignas (float)       int32_t j;
extern alignas (float)       int32_t k;
extern alignas (float)       int32_t l;
extern alignas (float32_t)   int32_t m;

/* file1.c */

alignas (16)          int32_t a; /* Compliant - same explicit alignment */
alignas (16)          int32_t b; /* Non-compliant - not consistently explicit */
alignas (16)          int32_t c; /* Non-compliant - not consistently explicit */
                      int32_t d; /* Compliant - not manually aligned */
                      int32_t e; /* Non-compliant - constraint violation */
alignas (16)          int32_t f; /* Non-compliant because of file2.c */
alignas (16)          int32_t g; /* Non-compliant, and undefined because of file2.c */

extern alignas (16) int32_t h; /* Non-compliant with file2.c */

alignas (float)       int32_t i; /* Compliant - same type used */
alignas (double)      int32_t j; /* Non-compliant - different type, and therefore
                                may be a constraint violation */
alignas (4)           int32_t k; /* Non-compliant - regardless of the size of float */
alignas (float32_t)   int32_t l; /* Non-compliant - potentially a different type on
                                a different platform */
alignas (float32_t)   int32_t m; /* Compliant - same type used by name */

/* file2.c */
extern                      int32_t f; /* Non-compliant - not consistent with
                                either file1.c or header.h */

extern alignas (8)    int32_t g; /* Non-compliant - undefined behaviour because of
                                inconsistency with file1.c */

extern alignas (8)    int32_t h; /* Non-compliant - not consistent with file1.c
                                and undefined behaviour */

```

See also

Rule 8.6, Rule 8.16, Rule 8.17

Rule 8.16 The *alignment specification* of zero should not appear in an object declaration

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C11

Amplification

This rule applies to any *integer constant expression* operand to `_Alignas` that evaluates to zero.

Rationale

If the alignment of an object is important, it should be specified explicitly.

If configuration settings or platform implementation details are intended to change the alignment of an object to conditionally disable explicit alignment, this should be abstracted by the preprocessor.

Example

```

int32_t a; /* Compliant: no alignment specification */
alignas (16) int32_t b; /* Compliant: explicit non-zero alignment specification */
alignas (0) int32_t c; /* Non-compliant: zero-alignment specification */

/* Non-compliant on platforms where sizeof (int) == sizeof (long) */
alignas (sizeof (long) - sizeof (int)) int32_t d;

```

When the alignment is not important, the configuration can remove it entirely:

```

#if REQUIRED_ALIGNMENT > 0
#define ALIGNED alignas (REQUIRED_ALIGNMENT)
#else
#define ALIGNED /**/
#endif

```

See also

Rule 8.15, Rule 8.17

Rule 8.17 At most one explicit *alignment specifier* should appear in an object declaration

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C11

Rationale

If the alignment of an object is important, it should be specified explicitly.

Because an *alignment specifier* only places a minimum requirement on the actual alignment of an object, C permits a declaration to contain multiple *alignment specifiers*, with the strictest imposing the final requirement.

If separate conditions require different minimum permitted alignments for an object, they should be combined explicitly by the expression controlling the specifier. Otherwise, the conflicting *alignment specifiers* risk obscuring the intent of the declaration from a human reviewer.

Example

```

int32_t a; /* Compliant - no alignment specifier */
alignas (16) int32_t b; /* Compliant - one alignment specifier */
alignas (16) alignas (8) int32_t c; /* Non-compliant - two alignment specifiers */
alignas (16) alignas (0) int32_t d; /* Non-compliant - also violates Rule 8.16 */

```

The following example shows a way of generating conditional alignment:

```
#define SML_ALIGN 16
#define BIG_ALIGN 32
alignas(MAX(SML_ALIGN, BIG_ALIGN)) int32_t e; /* Compliant */
```

See also

Rule 8.15, Rule 8.16

8.9 Initialization

Rule 9.1 The value of an object with automatic storage duration shall not be read before it has been set

C90 [Undefined 41], C99 [Undefined 10, 17], C11 [Undefined 11, 19]

Category Mandatory

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

For the purposes of this rule, an array element or structure member shall be considered as a discrete object.

This rule does not apply to *_Atomic* qualified objects, which are covered by Rule 9.7.

Rationale

According to the C Standard, objects with static storage duration are automatically initialized to zero unless initialized explicitly. Objects with automatic storage duration are not automatically initialized and can therefore have indeterminate values.

Note: it is sometimes possible for the explicit initialization of an automatic object to be ignored. This will happen when a jump to a label using a *goto* or *switch* statement “bypasses” the declaration of the object; the object will be declared as expected but any explicit initialization will be ignored.

Example

```
void f ( bool_t b, uint16_t *p )
{
    if ( b )
    {
        *p = 3U;
    }
}

void g ( void )
{
    uint16_t u;

    f ( false, &u );

    if ( u == 3U )
    {
        /* Non-compliant - u has not been assigned a value */
    }
}
```

In the following non-compliant C99 example, the `goto` statement jumps past the initialization of `x`.

Note: This example is also non-compliant with Rule 15.1.

```
{
    goto L1;

    uint16_t x = 10u;

L1:
    x = x + 1u;    /* Non-compliant - x has not been assigned a value */
}
```

See also

Rule 15.1, Rule 9.7, Rule 15.3

Rule 9.2 The initializer for an aggregate or union shall be enclosed in braces

C90 [Undefined 42], C99 [Undefined 76, 77], C11 [Undefined 82, 83]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule applies to initializers for both objects and subobjects.

An initializer of the form `{ 0 }`, which sets all values to 0, may be used to initialize subobjects without nested braces.

Note: this rule does not itself require explicit initialization of objects or subobjects.

Rationale

Using braces to indicate initialization of subobjects improves the clarity of code and forces programmers to consider the initialization of elements in complex data structures such as multi-dimensional arrays or arrays of structures.

Exception

1. An array may be initialized using a string literal.
2. An automatic structure or union may be initialized using an expression with compatible structure or union type.
3. A designated initializer may be used to initialize part of a subobject.

Example

The following three initializations, which are permitted by the C Standard, are equivalent. The first form is not permitted by this rule because it does not use braces to show subarray initialization explicitly.

```
int16_t y[ 3 ][ 2 ] = { 1, 2, 0, 0, 5, 6 };          /* Non-compliant */
int16_t y[ 3 ][ 2 ] = { { 1, 2 }, { 0 }, { 5, 6 } }; /* Compliant */
int16_t y[ 3 ][ 2 ] = { { 1, 2 }, { 0, 0 }, { 5, 6 } }; /* Compliant */
```

In the following example, the initialization of `z1` is compliant by virtue of Exception 3 because a designated initializer is used to initialize the subobject `z1[1]`. The initialization of `z2` is also compliant for the same reason. The initialization of `z3` is non-compliant because part of the subobject `z3[1]` is initialized with a positional initializer but is not enclosed in braces. The initialization of `z4` is compliant because a designated initializer is used to initialize the subobject `z4[0]` and the initializer for subobject `z4[1]` is brace-enclosed.

```
int16_t z1[ 2 ][ 2 ] = { { 0 }, [ 1 ][ 1 ] = 1 };          /* Compliant */
int16_t z2[ 2 ][ 2 ] = { { 0 },
                        [ 1 ][ 1 ] = 1, [ 1 ][ 0 ] = 0
                        };                               /* Compliant */
int16_t z3[ 2 ][ 2 ] = { { 0 }, [ 1 ][ 0 ] = 0, 1 };    /* Non-compliant */
int16_t z4[ 2 ][ 2 ] = { [ 0 ][ 1 ] = 0, { 0, 1 } };   /* Compliant */
```

The first line in the following example initializes 3 subarrays without using nested braces. The second and third lines show equivalent ways to write the same initializer.

```
float32_t a[ 3 ][ 2 ] = { 0 };                          /* Compliant */
float32_t a[ 3 ][ 2 ] = { { 0 }, { 0 }, { 0 } };        /* Compliant */
float32_t a[ 3 ][ 2 ] = { { 0.0f, 0.0f },
                          { 0.0f, 0.0f },
                          { 0.0f, 0.0f }
                          };                             /* Compliant */

union u1 {
    int16_t i;
    float32_t f;
} u = { 0 };                                           /* Compliant */

struct s1 {
    uint16_t len;
    char buf[ 8 ];
} s[ 3 ] = {
    { 5u, { 'a', 'b', 'c', 'd', 'e', '\0', '\0', '\0' } },
    { 2u, { 0 } },
    { .len = 0u }
};                                                     /* Compliant - buf initialized implicitly */
                                                     /* Compliant - s[] fully initialized */
```

See also

Rule 9.6

Rule 9.3 Arrays shall not be partially initialized

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

If any element of an array object or subobject is explicitly initialized, then the entire object or subobject shall be explicitly initialized.

Rationale

Providing an explicit initialization for each element of an array makes it clear that every element has been considered.

Exception

1. An initializer of the form `{ 0 }` may be used to explicitly initialize all elements of an array object or subobject.
2. An array whose initializer consists **only** of designated initializers may be used, for example to perform a sparse initialization.
3. An array initialized using a string literal does not need an initializer for every element.

Example

```
/* Compliant */
int32_t x[ 3 ] = { 0, 1, 2 };

/* Non-compliant - y[ 2 ] is implicitly initialized */
int32_t y[ 3 ] = { 0, 1 };

/* Non-compliant - t[ 0 ] and t[ 3 ] are implicitly initialized */
float32_t t[ 4 ] = { [ 1 ] = 1.0f, 2.0f };

/* Compliant - designated initializers for sparse matrix */
float32_t z[ 50 ] = { [ 1 ] = 1.0f, [ 25 ] = 2.0f };
```

In the following compliant example, each element of the array `arr` is initialized:

```
float32_t arr[ 3 ][ 2 ] =
{
  { 0.0f, 0.0f },
  { PI / 4.0f, -PI / 4.0f },
  { 0 } /* initializes all elements of array subobject arr[ 2 ] */
};
```

In the following example, array elements 6 to 9 are implicitly initialized to `'\0'`:

```
char h[ 10 ] = "Hello"; /* Compliant by Exception 3 */
```

Rule 9.4 An element of an object shall not be initialized more than once

Category Required

Analysis Decidable, Single Translation Unit

Applies to C99, C11

Amplification

This rule applies to initializers for both objects and subobjects.

An aggregate initializer shall not contain two designators that refer to the same sub-object. An aggregate initializer shall not allow the *current object* to implicitly initialize an element that has been initialized previously in the initializer list.

Rationale

The provision of *designated initializers* allows the naming of the components of an aggregate (structure or array) or of a union to be initialized within an initializer list and allows the object's elements to be initialized in any order by specifying the array indices or structure member names they apply to (elements having no initialization value assume the default for uninitialized objects).

A designator can specify elements to be initialized in a different syntactic sequence from their order within the object layout. An initializer without a designator will always initialize the *next subobject* within the object layout.

Care is required when using *designated initializers* since the initialization of object elements can be inadvertently repeated. The C Standard specifies that the value produced by the syntactically-last initializer referring to an element in the list is used, overriding any preceding initializers for that element. The Standard leaves unspecified whether overridden initializers are evaluated, and therefore whether or not any *side effects* in the initializing expressions occur or not. This is not listed in Annex J of the C Standard.

In order to allow sparse arrays and structures, it is acceptable to only initialize those which are necessary to the application.

Example

Array initialization:

```
/*
 * Required behaviour using positional initialization
 * Compliant - a1 is -5, -4, -3, -2, -1
 */
int16_t a1[ 5 ] = { -5, -4, -3, -2, -1 };

/*
 * Similar behaviour using designated initializers
 * Compliant - a2 is -5, -4, -3, -2, -1
 */
int16_t a2[ 5 ] = { [ 0 ] = -5, [ 1 ] = -4, [ 2 ] = -3,
                  [ 3 ] = -2, [ 4 ] = -1 };

/*
 * Repeated designated initializer element values overwrite earlier ones
 * Non-compliant - a3 is -5, -4, -2, 0, -1
 */
int16_t a3[ 5 ] = { [ 0 ] = -5, [ 1 ] = -4, [ 2 ] = -3,
                  [ 2 ] = -2, [ 4 ] = -1 };
```

In the following non-compliant example, it is unspecified whether the *side effect* occurs or not:

```
uint16_t *p;

void f ( void )
{
    uint16_t a[ 2 ] = { [ 0 ] = *p++, [ 0 ] = 1 };
}
```

Structure initialization:

```
struct mystruct
{
    int32_t a;
    int32_t b;
    int32_t c;
    int32_t d;
};

/*
 * Required behaviour using positional initialization
 * Compliant - s1 is 100, -1, 42, 999
 */
struct mystruct s1 = { 100, -1, 42, 999 };

/*
 * Similar behaviour using designated initializers
 * Compliant - s2 is 100, -1, 42, 999
 */
struct mystruct s2 = { .a = 100, .b = -1, .c = 42, .d = 999 };
```

```

/*
 * Repeated designated initializer element values overwrite earlier ones
 * Non-compliant - s3 is 42, -1, 0, 999
 */
struct mystruct s3 = { .a = 100, .b = -1, .a = 42, .d = 999 };

/*
 * Positional initializer element values can overwrite earlier ones
 * if preceded by a designated element out of sequence
 * Non-compliant - s4 is 1, 4, 3, 0
 */
struct mystruct s4 = { .b = 2, .c = 3, .a = 1, /* b */ 4 };

```

See also

Rule 9.6

Rule 9.5 Where designated initializers are used to initialize an array object the size of the array shall be specified explicitly

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Amplification

The rule applies equally to an array subobject that is a flexible array member.

Rationale

If the size of an array is not specified explicitly, it is determined by the highest index of any of the elements that are initialized. When using designated initializers it may not always be clear which initializer has the highest index, especially when the initializer contains a large number of elements.

To make the intent clear, the array size shall be declared explicitly. This provides some protection if, during development of the program, the indices of the initialized elements are changed as it is a *constraint* violation to initialize an element outside the bounds of an array.

Example

```

/* Non-compliant - probably unintentional to have single element */
int a1[ ] = { [ 0 ] = 1 };

/* Compliant */
int a2[ 10 ] = { [ 0 ] = 1 };

```

Rule 9.6 An initializer using chained designators shall not contain initializers without designators

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Amplification

A chained designator is a *designator list* that has more than one item, thus specifying an element of a sub-object within the *current object*.

If an aggregate initializer uses designators to specify elements, and any designator in the initializer is chained, every initializer in the entire containing initializer list shall specify an element explicitly using a designator.

This rule applies to initializers for both objects and sub-objects.

Rationale

Using chained designators for selective sub-object designation can make the intent of the initializer clear for some constructs such as sparse matrices. However, combining chained designators with positional initialization is extremely unclear — a human reader cannot easily tell whether the intended *next object* is within the sub-object, or within the same object level from which the designator started lookup. The syntactic brace structure of the initializer list may also no longer match the depth of the selected element, adding to the confusion.

Exception

A braced sub-object initializer may omit designators to specify elements if it does not contain any chained designators, and no chained designators in the containing initializer list specify an element inside it as the *current object*.

Example

```
struct S
{
    int x;
    int y;
};

struct T
{
    int    w;
    struct S s;
    int    z;
};

/* Non-compliant - chained designators and implicit positional initializers mixed */
struct T tt = {
    1,
    .s.x = 2, /* To a human reader, this looks like .z is being initialized */
    3,       /* tt is actually initialized as { 1, { 2, 3 }, 0 } */
};          /* This also violates Rule 9.2 */
```

```

/* Compliant - allow the y dimension to implicitly initialize to zero */
struct S aa[5] = {
    [0].x = 1,
    [1].x = 2,
    [2].x = 3,
    [3].x = 4,
    [4].x = 5,
};

/* Compliant - the initializer for [1] is not chained, but is explicit */
struct S ab[2] = {
    [0].x = 1,
    [1] = { 2, 3 }, /* Compliant by exception: */
}; /* the positional initializers are inside a braced sub-object */

```

See also

Rule 9.2, Rule 9.4

Rule 9.7 Atomic objects shall be appropriately initialized before being accessed

C11 [Undefined 5, *]

Category Mandatory

Analysis Undecidable, System

Applies to C11

Amplification

Initialization of atomic objects shall be completed before accessing them:

- For objects that do not have static storage duration, initialization shall be included in their declaration using the assignment operator = or using the Standard Library function *atomic_init()*, before any other access.
- For objects of static storage duration, the default initialization is sufficient.

Rationale

An atomic object is to be initialized before it is accessed. Concurrent access to the object being initialized, even via an atomic operation, constitutes a data race.

The *atomic_init()* function initializes atomic objects, including any additional state that the implementation might need to carry for the atomic object. However, it does not avoid data races.

Because of the potential initialization of the implementation state, *atomic_init()* cannot be replaced by other access functions, e.g. *atomic_store()*. Initialization of atomic objects inside of threads would impose constraints on thread ordering which are hard to ensure or verify. An explicit protection, e.g. by use of a mutex, would make atomicity unnecessary.

Example

```

_Atomic int32_t g_ai1;          /* Compliant    - default initialization    */
void main( void )
{
  _Atomic int32_t ai1 = 22;     /* Compliant    - directly initialized    */
  _Atomic int32_t ai2;
  ai2 = 777;                   /* Non-compliant - not initialized by atomic_init */

  _Atomic int32_t ai3;
  atomic_init( &ai3, 333);     /* Compliant    - Initialized by atomic_init    */

  /* ----- */

  _Atomic int32_t ai4;
  thrd_create( &id1, t1, &ai4);

  atomic_init( &ai4, 666);     /* Non-compliant - Initialized after user-thread
                               T1 is created */

  thrd_join ( id1, NULL);
}

int32_t t1( t1_paramtype *ptr )
{
  /* accesses g_ai1, ai1, ai2, ai3, ai4 */
}

```

See also

Dir 5.1, Rule 1.5, Rule 9.1, Rule 12.6

8.10 The *essential type model*

8.10.1 Rationale

The rules in this section collectively define the *essential type model* and restrict the C type system so as to:

1. Support a stronger system of type-checking;
2. Provide a rational basis for defining rules to control the use of implicit and explicit type conversions;
3. Promote portable coding practices;
4. Address some of the type conversion anomalies found within ISO C.

The *essential type model* does this by allocating an *essential type* to those objects and expressions which ISO C considers to be of arithmetic type. For example, adding an *int* to a *char* gives a result having *essentially character* type rather than the *int* type that is actually produced by integer promotion.

The full rationale behind the *essential type model* is given in Appendix C with Appendix D providing a comprehensive definition of the *essential type* of any arithmetic expression.

The rules in this section do not apply to expressions with a pointer type, unless otherwise specified.

8.10.2 Essential type

The *essential type* of an object or expression is defined by its *essential type category* and size.

The *essential type category* of an expression reflects its underlying behaviour and may be:

- *Essentially Boolean*;
- *Essentially character*;
- *Essentially enum*;
- *Essentially signed*;
- *Essentially unsigned*;
- *Essentially floating*, which may be either:
 - *Essentially real floating*, or
 - *Essentially complex floating*.

Notes:

1. For the purposes of this *essential type model*, *essentially real floating* and *essentially complex floating* are considered to be distinct *essential type categories*.
2. Each enumerated type is a unique *essentially enum type* identified as *enum<i>*. This allows different enumerated types to be handled as distinct types, which supports a stronger system of type-checking. One exception is the use of an enumerated type to define a Boolean value in C90. Such types are considered to have *essentially Boolean* type. Another exception is the use of *anonymous enumerations* as defined in Appendix D. *Anonymous enumerations* are a way of defining a set of related constant integers and are considered to have an *essentially signed* type.

When comparing two types of the same type category, the terms *wider* and *narrower* are used to describe their relative sizes as measured in bytes. Two different types are sometimes implemented with the same size.

The following table shows how the standard integer types map on to *essential type categories*:

<i>Essential type category</i>				
Boolean	character	signed	unsigned	enum<i>
_Bool	char	signed char signed short signed int signed long signed long long	unsigned char unsigned short unsigned int unsigned long unsigned long long	named enum

<i>Essential type category</i>	
floating	
real floating	complex floating
float double long double	float _Complex double _Complex long double _Complex

Note: Implementations of C99 and later may provide *extended integer types*, each of which would be allocated a location appropriate to its rank and signedness.

The restrictions enforced by the rules in this section also apply to the compound assignment operators (e.g. +=) as these are equivalent to assigning the result obtained from the use of one of the arithmetic, bitwise or shift operators. For example:

```
u8a += u8b + 1U;
```

is equivalent to:

```
u8a = u8a + ( u8b + 1U );
```

Rule 10.1 Operands shall not be of an inappropriate *essential type*

C90 [Unspecified 23; Implementation 14, 17, 19, 32]

C99 [Undefined 13, 48, 49; Implementation J.3.4(2, 5), J.3.5(5), J.3.9(6)]

C11 [Undefined 15, 51, 52; Implementation J.3.4(2, 5), J.3.5(5), J.3.9(7)]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

In the following table a number within a cell indicates where a restriction applies to the use of an *essential type* as an operand to an operator. These numbers correspond to paragraphs in the Rationale section below and indicate why each restriction is imposed.

Operator	Operand	Essential type category of arithmetic operand						
		Boolean	character	enum	signed	unsigned	floating	
							real	complex
[]	integer	3	4				1	9
+ (unary)		3	4	5				
- (unary)		3	4	5		8		
+ -	either	3		5				
++ --		3		5				9
* /	either	3	4	5				
%	either	3	4	5			1	9
< > <= >=	either	3						9
== !=	either						10	10
! &&	any		2	2	2	2	2	2
<< >>	left	3	4	5, 6	6		1	9
<< >>	right	3	4	7	7		1	9
~ & ^	any	3	4	5, 6	6		1	9
?:	1st		2	2	2	2	2	2
?:	2nd and 3rd							

In addition, the rule prohibits the use of logical operators (! && ||) on an operand with pointer type.

Other rules place further restrictions on the combination of *essential types* that may be used within an expression.

Rationale

1. The use of an expression of *essentially floating type* for these operands is a *constraint violation*.
2. An expression of *essentially Boolean type* should always be used where an operand is interpreted as a Boolean value.
3. An operand of *essentially Boolean type* should not be used where an operand is interpreted as a numeric value.
4. An operand of *essentially character type* should not be used where an operand is interpreted as a numeric value. The numeric values of character data are implementation-defined.
5. An operand of *essentially enum type* should not be used in an arithmetic operation because an *enum* object uses an implementation-defined integer type. An operation involving an *enum* object may therefore yield a result with an unexpected type. Note that an enumeration constant from an *anonymous enum* has *essentially signed type*.
6. Shift and bitwise operations should only be performed on operands of *essentially unsigned type*. The numeric value resulting from their use on *essentially signed types* may be undefined or implementation-defined.
7. The right hand operand of a shift operator should be of *essentially unsigned type* to ensure that undefined behaviour does not result from a negative shift.
8. An operand of *essentially unsigned type* should not be used as the operand to the unary minus operator, as the signedness of the result is determined by the implemented size of *int*.
9. The use of an expression of *essentially complex floating type* for these operands is a constraint violation.
10. The inherent nature of floating-point types is such that comparisons of equality will often not evaluate to true even when they are expected to. In addition the behaviour of such a comparison cannot be predicted before execution, and may well vary from one implementation to another. Deterministic floating-point comparisons should take into account the floating-point granularity (`FLT_EPSILON`) and the magnitude of the numbers being compared.

Exception

1. A non-negative *integer constant expression* of *essentially signed type* may be used as the right hand operand to a shift operator.
2. Comparison (for equality or inequality) of *essentially real floating* or *essentially complex floating* expressions with the constant literal value of zero or with the macro `INFINITY` or `-INFINITY` is permitted.

Example

```
enum enuma { a1, a2, a3 } ena, enb;    /* Essentially enum<enuma>          */
enum { K1 = 1, K2 = 2 };              /* Essentially signed                */
```

The following examples are non-compliant. The comments refer to the numbered rationale item that results in the non-compliance.

```
f32a & 2U          /* Rationale 1 - constraint violation */
f32a << 2          /* Rationale 1 - constraint violation */
```

```

cha && bla      /* Rationale 2 - char type used as a Boolean value */
ena ? a1 : a2   /* Rationale 2 - enum type used as a Boolean value */
s8a && bla      /* Rationale 2 - signed type used as a Boolean value */
u8a ? a1 : a2   /* Rationale 2 - unsigned type used as a Boolean value */
f32a && bla     /* Rationale 2 - floating type used as a Boolean value */

bla * blb      /* Rationale 3 - Boolean used as a numeric value */
bla > blb      /* Rationale 3 - Boolean used as a numeric value */

cha & chb      /* Rationale 4 - char type used as a numeric value */
cha << 1       /* Rationale 4 - char type used as a numeric value */

ena--          /* Rationale 5 - enum type used in arithmetic operation */
ena * a1       /* Rationale 5 - enum type used in arithmetic operation */

s8a & 2        /* Rationale 6 - bitwise operation on signed type */
50 << 3U       /* Rationale 6 - shift operation on signed type */

u8a << s8a     /* Rationale 7 - shift magnitude uses signed type */
u8a << -1      /* Rationale 7 - shift magnitude uses signed type */

-u8a          /* Rationale 8 - unary minus on unsigned type */

```

The following example is non-compliant with this rule and also violates Rule 10.3:

```
ena += a1      /* Rationale 5 - enum type used in arithmetic operation */
```

The following examples are compliant:

```

bla && blb
bla ? u8a : u8b

cha - chb
cha > chb

ena > a1
K1 * s8a      /* Compliant as K1 from anonymous enum */

s8a + s16b
-( s8a ) * s8b
s8a > 0
--s16b

u8a + u16b
u8a & 2U

u8a > 0U
u8a << 2U
u8a << 1      /* Compliant by Exception 1 */

f32a + f32b
f32a > 0.0

```

The following examples demonstrate the scope of Rationale 10:

```

float32_t f32w = -1.0f;
float32_t f32x = 0.2f;
float32_t f32y = 0.9f;
float32_t f32z = 0.0f;

if ( ( f32w + f32x ) == f32y ) /* Non-compliant */
if ( f32w != f32y )          /* Non-compliant */
if ( f32w <= f32y )          /* Compliant */

if ( f32w != 0.0f )          /* Compliant - Exception 2 */
if ( f32w != f32z )          /* Non-compliant - not a literal constant */

```

The following is compliant with this rule but violates Rule 10.2

```
cha + chb
```

See also

Dir 4.15, Rule 10.2

Rule 10.2 Expressions of *essentially character type* shall not be used inappropriately in addition and subtraction operations

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The appropriate uses are:

1. For the + operator, one operand shall have *essentially character type* and the other shall have *essentially signed type* or *essentially unsigned type* having a rank lower than or equal to that of *int*; the result of the operation has *essentially character type*.
2. For the - operator, the first operand shall have *essentially character type* and the second shall have *essentially signed type* or *essentially unsigned type* or *essentially character type*, and a rank lower than or equal to that of *int*; if both operands have *essentially character type* then the result has the *standard type* (usually *int* in this case) else the result has *essentially character type*.

Rationale

Expressions with *essentially character type* (character data) shall not be used arithmetically as the data does not represent numeric values.

The uses above are permitted as they allow potentially reasonable manipulation of character data. For example:

- Subtraction of two operands with *essentially character type* might be used to convert between digits in the range '0' to '9' and the corresponding ordinal value;
- Addition of an *essentially character type* and an *essentially unsigned type* might be used to convert an ordinal value to the corresponding digit in the range '0' to '9';
- Subtraction of an *essentially unsigned type* from an *essentially character type* might be used to convert a character from lowercase to uppercase.

Example

The following examples are compliant:

```
'0' + u8a /* Convert u8a to digit */
s8a + '0' /* Convert s8a to digit */
cha - '0' /* Convert cha to ordinal */
'0' - s8a /* Convert -s8a to digit */
```

The following examples are non-compliant:

```
s16a - 'a'
'0' + f32a
cha + ':'
cha - ena
```

See also

Rule 10.1

Rule 10.3 The value of an expression shall not be assigned to an object with a narrower *essential type* or of a different *essential type category*

C90 [Undefined 15; Implementation 16]

C99 [Undefined 15, 16; Implementation J.3.5(4)]

C11 [Undefined 17, 18; Implementation J.3.5(4)]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

The following operations are covered by this rule:

1. *Assignment* as defined in the Glossary;
2. The conversion of the *constant expression* in a *switch* statement's *case* label to the *essential type* of the controlling expression.

Rationale

The C language allows the programmer considerable freedom and will permit assignments between different arithmetic types to be performed automatically. However, the use of these implicit conversions can lead to unintended results, with the potential for loss of value, sign or precision. Further details of concerns with the C type system can be found in Appendix C.

The use of stronger typing, as enforced by the MISRA *essential type model*, reduces the likelihood of these problems occurring.

Exception

1. An *essentially signed integer constant expression*, with a rank no greater than signed int, may be assigned to an object of *essentially unsigned type* if its value can be represented in that type.
2. The initializer { 0 } may be used to initialize an aggregate or union type.
3. A *switch* statement's *case* label that is a non-negative *integer constant expression* of *essentially signed type* is permitted when the controlling expression is of *essentially unsigned type* and the value can be represented in that type.
4. An *essentially real floating* expression may be assigned to an object of *essentially complex floating type* provided that its *corresponding real type* is not narrower than the type of the expression.

Example

```
enum enuma { A1, A2, A3 } ena;
enum enumb { B1, B2, B3 } enb;
enum      { K1=1, K2=128 };
```

The following are compliant:

```
uint8_t u8a = 0;           /* By exception */
bool_t  flag = ( bool_t ) 0;
bool_t  set  = true;      /* true is essentially Boolean */
bool_t  get  = ( u8b > u8c );

ena  = A1;
s8a  = K1;                /* Constant value fits */
u8a  = 2;                 /* By exception */
u8a  = 2 * 24;           /* By exception */
cha += 1;                 /* cha = cha + 1 assigns character to character */

pu8a = pu8b;             /* Same essential type */
u8a  = u8b + u8c + u8d;  /* Same essential type */
u8a  = ( uint8_t ) s8a;  /* Cast gives same essential type */

u32a = u16a;             /* Assignment to a wider essential type */
u32a = 2U + 125U;       /* Assignment to a wider essential type */
use_uint16 ( u8a );     /* Assignment to a wider essential type */
use_uint16 ( u8a + u16b ); /* Assignment to same essential type */

cf32a = f32a;           /* By exception 4 */
cf64a = f64a;           /* By exception 4 */
```

The following are non-compliant as they have different *essential type categories*:

```
uint8_t u8a = 1.0f;      /* unsigned and floating */
bool_t  bla = 0;        /* boolean and signed */
char    = 7;           /* character and signed */
u8a    = 'a';          /* unsigned and character */
u8b    = 1 - 2;        /* unsigned and signed */
u8c    += 'a';         /* u8c = u8c + 'a' assigns character to unsigned */
use_uint32 ( s32a );    /* signed and unsigned */

f32a   = cf32a;        /* real floating and complex floating */
f64a   = cf64a;        /* real floating and complex floating */
```

The following are non-compliant as they contain assignments to a narrower *essential type*:

```
s8a    = K2;           /* Constant value does not fit */
u16a   = u32a;        /* uint32_t to uint16_t */

use_uint16 ( u32a );  /* uint32_t to uint16_t */

uint8_t fool ( uint16_t x )
{
    return x;         /* uint16_t to uint8_t */
}

cf32a  = f64a;        /* complex floating and real floating */
```

See also

Rule 10.4, Rule 10.5, Rule 10.6

Rule 10.4 Both operands of an operator in which the *usual arithmetic conversions* are performed shall have the same *essential type category*

C90 [Implementation 21], C99 [Implementation J.3.6(4)], C11 [Implementation J.3.6(5)]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

This rule applies to operators that are described in *usual arithmetic conversions* section in the C Standard. This includes all the binary operators, excluding the shift, logical &&, logical || and comma operators. In addition, the second and third operands of the ternary operator are covered by this rule.

Note: the increment and decrement operators are not covered by this rule.

Rationale

The C language allows the programmer considerable freedom and will permit conversions between different arithmetic types to be performed automatically. However, the use of these implicit conversions can lead to unintended results, with the potential for loss of value, sign or precision. Further details of concerns with the C type system can be found in Appendix C.

The use of stronger typing, as enforced by the MISRA *essential type model*, allows implicit conversions to be restricted to those that should then produce the answer expected by the developer.

Exception

The following are permitted to allow a common form of character manipulation to be used:

1. The binary + and += operators may have one operand with *essentially character* type and the other operand with an *essentially signed* or *essentially unsigned* type;
2. The binary – and -= operators may have a left-hand operand with *essentially character* type and a right-hand operand with an *essentially signed* or *essentially unsigned* type.

Operations involving mixed real and complex operands have well-defined semantics and are therefore permitted:

3. The operators covered by this rule may have one operand with *essentially real floating* type and the other operand with *essentially complex floating* type. In the case of the conditional operator, this exception applies to the second and third operands.

Example

```
enum enuma { A1, A2, A3 } ena;
enum enumb { B1, B2, B3 } enb;
```

The following are compliant as they have the same *essential type category*:

```
ena > A1
u8a + u16b
```

The following is compliant by exception 1:

```
cha += u8a
```

The following is compliant by exception 1, but violates Rule 10.3:

```
u8a += cha      /* unsigned and char          */
```

The following is compliant by exception 3.

```
cf32a += f32a; /* complex floating and real floating */
```

The following is non-compliant with this rule and also violates Rule 10.3:

```
s8a += u8a      /* signed and unsigned          */
```

The following are non-compliant:

```
u8b + 2         /* unsigned and signed          */
enb > A1        /* enum<enumb> and enum<enuma>  */
ena == enb      /* enum<enuma> and enum<enumb>  */
```

See also

Rule 10.3, Rule 10.7

Rule 10.5 The value of an expression should not be cast to an inappropriate *essential type*

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

The casts which should be avoided are shown in the following table, where values are cast (explicitly converted) to the *essential type category* of the first column.

Essential type category	from						
	<i>Boolean</i>	<i>character</i>	<i>enum</i>	<i>signed</i>	<i>unsigned</i>	<i>real floating</i>	<i>complex floating</i>
to							
<i>Boolean</i>	-	Avoid	Avoid	Avoid	Avoid	Avoid	Avoid
<i>character</i>	Avoid	-				Avoid	Avoid
<i>enum</i>	Avoid	Avoid	Avoid*	Avoid	Avoid	Avoid	Avoid
<i>signed</i>	Avoid			-			
<i>unsigned</i>	Avoid				-		
<i>real floating</i>	Avoid	Avoid				-	
<i>complex floating</i>	Avoid	Avoid					-

Note: an enumerated type may be cast to an enumerated type provided that the cast is to the same *essential enumerated* type. Such casts are redundant.

Casting from *void* to any other type is not permitted as it results in undefined behaviour. This is covered by Rule 1.3.

Rationale

An explicit cast may be introduced for legitimate functional reasons, for example:

- To change the type in which a subsequent arithmetic operation is performed;
- To truncate a value deliberately;
- To make a type conversion explicit in the interests of clarity.

However, some explicit casts are considered inappropriate:

- In C99 and later, the result of a cast or assignment to `_Bool` is always 0 or 1. This is not necessarily the case when casting to another type which is defined as *essentially Boolean*;
- A cast to an *essentially enum type* may result in a value that does not lie within the set of enumeration constants for that type;
- A cast from *essentially Boolean* to any other type is unlikely to be meaningful;
- Converting between floating and character types is not meaningful as there is no precise mapping between the two representations.

Exception

An *integer constant expression* with the value 0 or 1 and either *essentially signed* or *essentially unsigned* type may be cast to a type which is defined as *essentially Boolean*. This allows the implementation of Boolean models in C90.

Example

```
( bool_t ) false      /* Compliant - 'false' from stdbool.h is essentially Boolean */
( int32_t ) 3U        /* Compliant */
( bool_t ) 0          /* Compliant - by exception */
( bool_t ) 3U         /* Non-compliant */

( int32_t ) ena       /* Compliant */
( enum enuma ) 3     /* Non-compliant */
( char ) enc          /* Compliant */
```

See also

Rule 10.3, Rule 10.8

8.10.3 Composite operators and expressions

Some of the concerns mentioned in Appendix C can be avoided by restricting the implicit and explicit conversions that may be applied to non-trivial expressions. These include:

- The confusion about the type in which integer expressions are evaluated, as this depends on the type of the operands after any integer promotion. The type of the result of an arithmetic operation depends on the implemented size of *int*;
- The common misconception among programmers that the type in which a calculation is conducted is influenced by the type to which the result is assigned or cast. This false expectation may lead to unintended results.

In addition to the previous rules, the *essential type model* places further restrictions on expressions whose operands are *composite expressions*, as defined below.

The following are defined as *composite operators* in this document:

- Multiplicative (*, /, %)
- Additive (binary +, binary -)
- Bitwise (&, |, ^, ~)
- Shift (<<, >>)
- Conditional (:?) if either the second or third operand is a *composite expression*

A compound assignment is equivalent to an assignment of the result of its corresponding *composite operator*.

A *composite expression* is defined in this document as a non-constant *expression* which is the direct result of a *composite operator*.

Note:

- A parenthesized *composite expression* is also a *composite expression*;
- A unary + or unary - expression whose operand is a *composite expression* is also a *composite expression*.
- The results of the following operators are not *composite expressions*:
 - Assignment and compound assignment
 - Postfix and prefix increment and decrement
 - Cast
- A *constant expression* is not a *composite expression*.

Rule 10.6 The value of a *composite expression* shall not be assigned to an object with wider *essential type*

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

This rule covers the assigning operations described in Rule 10.3.

Rationale

The rationale is described in the introduction on *composite operators and expressions* (see Section 8.10.3).

Example

The following are compliant:

```
u16c = u16a + u16b;           /* Same essential type      */
u32a = ( uint32_t ) u16a + u16b; /* Cast causes addition in uint32_t */
```

The following are non-compliant:

```
u32a = u16a + u16b;           /* Implicit conversion on assignment */
use_uint32 ( u16a + u16b );   /* Implicit conversion of fn argument */
```

See also

Rule 10.3, Rule 10.7, Section 8.10.3

Rule 10.7 If a *composite expression* is used as one operand of an operator in which the *usual arithmetic conversions* are performed then the other operand shall not have wider *essential type*

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

The rationale is described in the introduction on *composite operators and expressions* (see Section 8.10.3).

Restricting implicit conversions on *composite expressions* means that sequences of arithmetic operations within an expression must be conducted in the same *essential type*, with the exception that *essentially real floating* and *essentially complex floating* operations may coexist. This reduces possible developer confusion.

Note: this does not imply that all operands in an expression are of the same *essential type*.

The expression `u32a + u16b + u16c` is compliant as both additions will notionally be performed in type `uint32_t`. In this case only *non-composite expressions* are implicitly converted.

The expression `(u16a + u16b) + u32c` is non-compliant as the left addition is notionally performed in type `uint16_t` and the right in type `uint32_t`, requiring an implicit conversion of the *composite expression* `u16a + u16b` to `uint32_t`.

Example

The following are compliant:

```
u32a * u16a + u16b           /* No composite conversion */
( u32a * u16a ) + u16b       /* No composite conversion */
u32a * ( ( uint32_t ) u16a + u16b ) /* Both operands of * have
                               * same essential type */
u32a += ( u32b + u16b )      /* No composite conversion */
```

The following are non-compliant:

```
u32a * ( u16a + u16b ) /* Implicit conversion of ( u16a + u16b ) */
u32a += ( u16a + u16b ) /* Implicit conversion of ( u16a + u16b ) */
```

See also

Rule 10.4, Rule 10.6, Section 8.10.3

Rule 10.8 The value of a *composite expression* shall not be cast to a different *essential type category* or a wider *essential type*

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

The rationale is described in the introduction on *composite operators and expressions* (see Section 8.10.3).

Casting to a wider type is not permitted as the result may vary between implementations. Consider the following:

```
( uint32_t ) ( u16a + u16b );
```

On a 16-bit machine the addition will be performed in 16 bits with the result wrapping modulo-2¹⁶ before it is cast to 32 bits. However, on a 32-bit machine the addition will take place in 32 bits and would preserve high-order bits that would have been lost on a 16-bit machine.

Casting to a narrower type with the same *essential type category* is acceptable as the explicit truncation of the result always leads to the same loss of information.

Exception

1. An *essentially real floating* expression may be cast to an *essentially complex floating* type providing that the corresponding real type is not wider than the type of the expression.
2. An *essentially real complex* expression may be cast to *essentially real floating* type providing that type is not wider than the corresponding real type of the expression.

Example

```
( uint16_t ) ( u32a + u32b ) /* Compliant */
( uint16_t ) ( s32a + s32b ) /* Non-compliant - different essential
                             type category */
( uint16_t ) s32a           /* Compliant - s32a is not composite */
( uint32_t ) ( u16a + u16b ) /* Non-compliant - cast to wider
                             essential type */
```

See also

Rule 10.5, Section 8.10.3

8.11 Pointer type conversions

Pointer types can be classified as follows:

- Pointer to object;
- Pointer to function;
- Pointer to incomplete;

- Pointer to *void*;
- A *null pointer constant*, i.e. the value 0, optionally cast to *void **.

The only conversions involving pointers that are permitted by the C Standard are:

- A conversion from a pointer type into *void*;
- A conversion from a pointer type into an arithmetic type;
- A conversion from an arithmetic type into a pointer type;
- A conversion from one pointer type into another pointer type.

Although permitted by the language *constraints*, conversion between pointers and any arithmetic types other than integer types is undefined.

The following permitted pointer conversions do **not** require an explicit cast:

- A conversion from a pointer type into *_Bool* (C99 and later);
- A conversion from a *null pointer constant* into a pointer type;
- A conversion from a pointer type into a compatible pointer type provided that the destination type has all the type qualifiers of the source type;
- A conversion between a pointer to an object or incomplete type and *void **, or qualified version thereof, provided that the destination type has all the type qualifiers of the source type.

Any implicit conversion that does not fall into this subset of pointer conversions either leads to undefined behaviour (C90) or violates a *constraint* (C99 and later).

The conversion between pointer types and integer types is implementation-defined.

Rule 11.1 Conversions shall not be performed between a pointer to a function and any other type

C90 [Undefined 24, 27–29], C99 [Undefined 21, 23, 39, 41], C11 [Undefined 24, 26, 41, 44]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

A pointer to a function shall only be converted into or from a pointer to a function with a compatible type.

Rationale

The conversion of a pointer to a function into or from any of:

- Pointer to object;
- Pointer to incomplete;
- *void **

results in undefined behaviour.

If a function is called by means of a pointer whose type is not compatible with the called function, the behaviour is undefined. Conversion of a pointer to a function into a pointer to a function with a different type is permitted by the C Standard. Conversion of an integer into a pointer to a function is also permitted by the C Standard. However, both are prohibited by this rule in order to avoid the undefined behaviour that would result from calling a function using an incompatible pointer type.

Exception

1. A *null pointer constant* may be converted into a pointer to a function;
2. A pointer to a function may be converted into *void*;
3. A function type may be implicitly converted into a pointer to that function type.

Note: exception 3 covers the implicit conversions that commonly occur when:

- A function is called directly, i.e. using a function identifier to denote the function to be called;
- A function is assigned to a function pointer.

Example

```
typedef void ( *fp16 ) ( int16_t n );
typedef void ( *fp32 ) ( int32_t n );

#include <stdlib.h>                                /* To obtain macro NULL      */
fp16 fp1 = NULL;                                  /* Compliant - exception 1   */
fp32 fp2 = ( fp32 ) fp1;                          /* Non-compliant - function
* pointer into different    */
* function pointer        */

if ( fp2 != NULL )                                /* Compliant - exception 1   */
{
}

fp16 fp3 = ( fp16 ) 0x8000;                        /* Non-compliant - integer into
* function pointer        */
fp16 fp4 = ( fp16 ) 1.0e6F;                        /* Non-compliant - float into
* function pointer        */
```

In the following example, the function call returns a pointer to a function type. Casting the return value into *void* is compliant with this rule.

```
typedef fp16 ( *pfp16 ) ( void );
pfp16 pfp1;

( void ) ( *pfp1 ( ) ); /* Compliant - exception 2 - cast function
* pointer into void    */
```

The following examples show compliant implicit conversions from a function type into a pointer to that function type.

```
extern void f ( int16_t n );

f ( 1 ); /* Compliant - exception 3 - implicit conversion
* of f into pointer to function */
fp16 fp5 = f; /* Compliant - exception 3 */
```

Rule 11.2 Conversions shall not be performed between a pointer to an incomplete type and any other type

C90 [Undefined 29], C99 [Undefined 21, 22, 41], C11 [Undefined 24, 25, 44]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

A pointer to an incomplete type shall not be converted into another type.

A conversion shall not be made into a pointer to incomplete type.

Although a pointer to *void* is also a pointer to an incomplete type, this rule does not apply to pointers to *void* as they are covered by Rule 11.5.

This rule applies to the unqualified types that are pointed to by the pointers.

Rationale

Conversion into or from a pointer to an incomplete type may result in a pointer that is not correctly aligned, resulting in undefined behaviour.

Conversion of a pointer to an incomplete type into or from a floating type always results in undefined behaviour.

Pointers to an incomplete type are sometimes used to hide the representation of an object. Converting a pointer to an incomplete type into a pointer to object would break this encapsulation.

Exception

1. A *null pointer constant* may be converted into a pointer to an incomplete type.
2. A pointer to an incomplete type may be converted into *void*.

Example

```

struct s;          /* Incomplete type          */
struct t;          /* A different incomplete type      */
struct s *sp;
struct t *tp;
int16_t *ip;

#include <stdlib.h> /* To obtain macro NULL            */

ip = ( int16_t * ) sp; /* Non-compliant                    */
sp = ( struct s * ) 1234; /* Non-compliant                    */
tp = ( struct t * ) sp; /* Non-compliant - casting pointer into a
                        * different incomplete type      */
sp = NULL;          /* Compliant - exception 1          */

struct s *f ( void );

( void ) f ( );     /* Compliant - exception 2          */

```

See also

Rule 11.5

Rule 11.3 A conversion shall not be performed between a pointer to object type and a pointer to a different object type

C90 [Undefined 20], C99 [Undefined 22, 34], C11 [Undefined 25, 37]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

This rule applies to the unqualified types that are pointed to by the pointers.

Rationale

Converting a pointer to object into a pointer to a different object may result in a pointer that is not correctly aligned, resulting in undefined behaviour.

Exception

It is permitted to convert a pointer to a non-qualified object type into a pointer to one of the object types *char*, *signed char* or *unsigned char*. The C Standard guarantees that pointers to these types can be used to access the individual bytes of an object.

Example

```
uint8_t *p1;
uint32_t *p2;

/* Non-compliant - possible incompatible alignment */
p2 = ( uint32_t * ) p1;

extern uint32_t read_value ( void );
extern void print ( uint32_t n );

void f ( void )
{
    uint32_t u    = read_value ( );
    uint16_t *hi_p = ( uint16_t * ) &u;    /* Non-compliant even though
                                           * probably correctly aligned    */

    *hi_p = 0;    /* Attempt to clear high 16-bits on big-endian machine */
    print ( u ); /* Line above may appear not to have been performed    */
}
```

The following example is compliant because the rule applies to the unqualified pointer types. It does not prevent type qualifiers from being added to the object type.

```
const short *p;
const volatile short *q;

q = ( const volatile short * ) p;    /* Compliant */
```

The following example is non-compliant because the unqualified pointer types are different, namely “pointer to *const*-qualified *int*” and “pointer to *int*”.

```
int * const * pcpi;
const int * const * pcpci;

pcpci = ( const int * const * ) pcpi;
```

See also

Rule 11.4, Rule 11.5, Rule 11.8, Rule 18.1

Rule 11.4 A conversion should not be performed between a pointer to object and an integer type

C90 [Undefined 20; Implementation 24]
 C99 [Undefined 21, 34; Implementation J.3.7(1)]
 C11 [Undefined 24, 37; Implementation J.3.7(1)]

Category Advisory
Analysis Decidable, Single Translation Unit
Applies to C90, C99, C11

Amplification

An object pointer should not be converted into an integer.

An integer should not be converted into an object pointer.

Rationale

Conversion of an integer into a pointer to object may result in a pointer that is not correctly aligned, resulting in undefined behaviour.

Conversion of a pointer to object into an integer may produce a value that cannot be represented in the chosen integer type resulting in undefined behaviour.

Note: the types `intptr_t` and `uintptr_t`, declared in `<stdint.h>` (C99 or later), are respectively signed and unsigned integer types capable of representing pointer values. Despite this, conversions between a pointer to object and these types is not permitted by this rule because their use does not avoid the undefined behaviour associated with misaligned pointers.

Casting between a pointer and an integer type should be avoided where possible, but may be necessary when addressing memory mapped registers or other hardware specific features. If casting between integers and pointers is used, care should be taken to ensure that any pointers produced do not give rise to the undefined behaviour discussed under Rule 11.3.

Exception

A *null pointer constant* that has integer type may be converted into a pointer to object.

Example

```
uint8_t *PORTA = ( uint8_t * ) 0x0002;    /* Non-compliant */
uint16_t *p;

int32_t  addr = ( int32_t ) &p;          /* Non-compliant */
uint8_t  *q   = ( uint8_t * ) addr;     /* Non-compliant */
bool_t   b    = ( bool_t ) p;           /* Non-compliant */

enum etag { A, B } e = ( enum etag ) p; /* Non-compliant */
```

See also

Rule 11.3, Rule 11.7, Rule 11.9

Rule 11.5 A conversion should not be performed from pointer to *void* into pointer to object

C90 [Undefined 20], C99 [Undefined 22, 34], C11 [Undefined 25, 37]

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

Conversion of a pointer to *void* into a pointer to object may result in a pointer that is not correctly aligned, resulting in undefined behaviour. It should be avoided where possible but may be necessary, for example when dealing with memory allocation functions. If conversion from a pointer to object into a pointer to *void* is used, care should be taken to ensure that any pointers produced do not give rise to the undefined behaviour discussed under Rule 11.3.

Exception

A *null pointer constant* that has type pointer to *void* may be converted into pointer to object.

Example

```
uint32_t *p32;
void      *p;
uint16_t *p16;

p  = p32;          /* Compliant - pointer to uint32_t into
                  * pointer to void                */
p16 = p;          /* Non-compliant                */
p  = ( void * ) p16; /* Compliant                */
p32 = ( uint32_t * ) p; /* Non-compliant                */
```

See also

Rule 11.2, Rule 11.3

Rule 11.6 A cast shall not be performed between pointer to *void* and an arithmetic type

C90 [Undefined 29; Implementation 24]
 C99 [Undefined 21, 41; Implementation J.3.7(1)]
 C11 [Undefined 24, 44; Implementation J.3.7(1)]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

Conversion of an integer into a pointer to *void* results in behaviour that is implementation-defined.

Conversion of a pointer to *void* into an integer may produce a value that cannot be represented in the chosen integer type resulting in undefined behaviour.

Conversion between any non-integer arithmetic type and pointer to *void* is undefined.

Exception

An *integer constant expression* with value 0 may be cast into pointer to *void*.

Example

```
void      *p;
uint32_t  u;

/* Non-compliant - implementation-defined */
p = ( void * ) 0x1234u;

/* Non-compliant - undefined                */
p = ( void * ) 1024.0f;

/* Non-compliant - implementation-defined */
u = ( uint32_t ) p;
```

Rule 11.7 A cast shall not be performed between pointer to object and a non-integer arithmetic type

C90 [Undefined 29; Implementation 24]
 C99 [Undefined 21, 41; Implementation J.3.7(1)]
 C11 [Undefined 24, 44; Implementation J.3.7(1)]

Category Required
Analysis Decidable, Single Translation Unit
Applies to C90, C99, C11

Amplification

For the purposes of this rule a non-integer arithmetic type means one of:

- *Essentially Boolean*;
- *Essentially character*;
- *Essentially enum*;
- *Essentially floating*.

Rationale

Conversion of an *essentially Boolean*, *essentially character* or *essentially enum* type into a pointer to object may result in a pointer that is not correctly aligned, resulting in undefined behaviour.

Conversion of a pointer to object into an *essentially Boolean*, *essentially character* or *essentially enum* type may produce a value that cannot be represented in the chosen integer type resulting in undefined behaviour.

Conversion of a pointer to object into or from an *essentially floating* type results in undefined behaviour.

Example

```
int16_t *p;
float32_t f;

f = ( float32_t ) p; /* Non-compliant */
p = ( int16_t * ) f; /* Non-compliant */
```

See also

Rule 11.4

Rule 11.8 A conversion shall not remove any *const*, *volatile* or *_Atomic* qualification from the type pointed to by a pointer

C90 [Undefined 12, 39, 40], C99 [Undefined 30, 61, 62], C11 [Undefined 33, 64, 65]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

Any attempt to remove the qualification associated with the addressed type is a violation of the principle of type qualification.

Note: the qualification referred to here is not the same as any qualification that may be applied to the pointer itself.

Some of the problems that might arise if a qualifier is removed from the addressed object are:

- Removing a *const* qualifier might circumvent the read-only status of an object and result in it being modified;
- Removing a *const* qualifier might result in an exception when the object is accessed;
- Removing a *volatile* qualifier might result in accesses to the object being optimized away.
- Removing an *_Atomic* qualifier might circumvent the lock status of an object and potentially result in memory corruption.

Note: removal of the *restrict* type qualifier (C99 and later) is benign.

Example

```

uint16_t      x;
uint16_t * const cpi = &x;    /* const pointer          */
uint16_t * const *pcpi;      /* pointer to const pointer */
uint16_t *    *ppi;
const uint16_t *pci;         /* pointer to const          */
volatile uint16_t *pvi;      /* pointer to volatile       */
uint16_t      *pi;

pi = cpi;                    /* Compliant - no conversion
                               no cast required          */
pi = (uint16_t *)pci;        /* Non-compliant            */
pi = (uint16_t *)pvi;        /* Non-compliant            */
ppi = (uint16_t * *)pcpi;    /* Non-compliant            */

typedef struct s {
    uint8_t a;
    uint8_t b;
} s_t;

int main( void )
{
    _Atomic s_t astr;
    s_t lstr = { 7U, 42U };
    s_t *sptr = &astr;      /* Non-compliant - removes _Atomic qualifier */
}

```

See also

Rule 11.3, Rule 11.10

Rule 11.9 The macro `NULL` shall be the only permitted form of integer *null pointer constant*

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

An *integer constant expression* with the value 0 shall be derived from expansion of the macro `NULL` if it appears in any of the following contexts:

- As the value being *assigned* to a pointer;
- As an operand of an `==` or `!=` operator whose other operand is a pointer;
- As the second operand of a `? :` operator whose third operand is a pointer;
- As the third operand of a `? :` operator whose second operand is a pointer.

Ignoring whitespace and any surrounding parentheses, any such *integer constant expression* shall represent the entire expansion of `NULL`.

Note: a *null pointer constant* of the form `(void *)0` is permitted, whether or not it was expanded from `NULL`.

Rationale

Using `NULL` rather than 0 makes it clear that a *null pointer constant* was intended.

Exception

The initializer `{ 0 }` may be used to initialize an aggregate or union type containing pointers.

Example

In the following example, the initialization of `p2` is compliant because the *integer constant expression* 0 does not appear in one of the contexts prohibited by this rule.

```
int32_t *p1 = 0;           /* Non-compliant */
int32_t *p2 = ( void * ) 0; /* Compliant   */
```

In the following example, the comparison between `p2` and `(void *)0` is compliant because the *integer constant expression* 0 appears as the operand of a cast and not in one of the contexts prohibited by this rule.

```
#define MY_NULL_1 0
#define MY_NULL_2 ( void * ) 0

if ( p1 == MY_NULL_1 ) /* Non-compliant */
{
}
if ( p2 == MY_NULL_2 ) /* Compliant   */
{
}
```

The following example is compliant because use of the macro `NULL` provided by the implementation is always permitted, even if it expands to an *integer constant expression* with value 0.

```
#include <stddef.h> /* To obtain macro NULL */
/* Could also use stdio.h, stdlib.h and others in hosted environments */

extern void f ( uint8_t *p );

/* Compliant for any conforming definition of NULL, such as:
 * 0
 * ( void * ) 0
 * ( ( 0 ) )
 * ( ( ( 1 - 1 ) ) )
 */
f ( NULL );
```

See also

Rule 11.4

Rule 11.10 The `_Atomic` qualifier shall not be applied to the incomplete type `void`

Category Required

Analysis Decidable, Single Translation Unit

Applies to C11

Rationale

The C Standard does not explicitly prohibit usage of the type `void` with the `_Atomic` qualifier. However, it does not provide a guarantee that a pointer to `_Atomic void` has any particular size or alignment requirement, so it cannot be assumed that is the same as for a pointer to an arbitrary type `_Atomic T`, and the behaviour of type conversion between them may be undefined.

Example

```
struct A {
    int32_t _Atomic x;
    int32_t _Atomic y;
};

void main (void)
{
    struct A a1 = { 6, 7 };

    void _Atomic * pav = &a1; /* Non-compliant */
    void _Atomic * pax = &a1.x; /* Non-compliant */
}
```

See also

Rule 11.8

8.12 Expressions

Rule 12.1 The precedence of operators within expressions should be made explicit

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification


The following table is used in the definition of this rule.

Description	Operator or Operand	Precedence
Primary	identifier, constant, string literal, (expression), generic selection	16 (high)
Postfix	[] () (function call) . -> ++ (post-increment) -- (post-decrement) () { } (compound literal)	15
Unary	++ (pre-increment) -- (pre-decrement) & * + - ~ ! <i>sizeof_Alignof defined</i> (preprocessor)	14
Cast	()	13
Multiplicative	* / %	12
Additive	+ -	11
Bitwise shift	<< >>	10
Relational	< > <= >=	9
Equality	== !=	8
Bitwise AND	&	7
Bitwise XOR	^	6
Bitwise OR		5
Logical AND	&&	4
Logical OR		3
Conditional	? :	2
Assignment	= *= /= %= += -= <<= >>= &= ^= =	1
Comma	,	0 (low)

The precedence used in this table are chosen to allow a concise description of the rule. They are not necessarily the same as those that might be encountered in other descriptions of operator precedence.

For the purposes of this rule, the precedence of an expression is the precedence of the element (operand or operator) at the root of the parse tree for that expression.

For example: the parse tree for the expression `a << b + c` can be represented as:



The element at the root of this parse tree is '<<' so the expression has precedence 10.

The following advice is given:

- The operand of the *sizeof* operator should be enclosed in parentheses;
- An expression whose precedence is in the range 2 to 12 should have parentheses around any operand that has both:
 - Precedence of less than 13, and
 - Precedence greater than the precedence of the expression.

Rationale

The C language has a relatively large number of operators and their relative precedence is not intuitive. This can lead less experienced programmers to make mistakes. Using parentheses to make operator precedence explicit removes the possibility that the programmer's expectations are incorrect. It also makes the original programmer's intention clear to reviewers or maintainers of the code.

It is recognized that overuse of parentheses can clutter the code and reduce its readability. This rule aims to achieve a compromise between code that is hard to understand because it contains either too many or too few parentheses.

Examples

The following example shows expressions with a unary or postfix operator whose operands are either *primary-expressions* or expressions whose top-level operators have precedence 15.

```
a[ i ]->n;      /* Compliant - no need to write ( a[ i ] )->n      */
*p++;          /* Compliant - no need to write *( p++ )          */
sizeof x + y;  /* Non-compliant - write either sizeof ( x ) + y          */
               * or sizeof ( x + y )                          */
```

The following example shows expressions containing operators at the same precedence level. All of these are compliant but, depending on the types of *a*, *b* and *c*, any expression with more than one operator may violate other rules.

```
a + b;
a + b + c;
( a + b ) + c;
a + ( b + c );
a + b - c + d;
( a + b ) - ( c + d );
```

The following example shows a variety of mixed-operator expressions:

```
/* Compliant - no need to write f ( ( a + b ), c ) */
x = f ( a + b, c );

/* Non-compliant
 * Operands of conditional operator (precedence 2) are:
 * == precedence 8 needs parentheses
 * a precedence 16 does not need parentheses
 * - precedence 11 needs parentheses
 */
x = a == b ? a : a - b;

/* Compliant */
x = ( a == b ) ? a : ( a - b );
```

```

/* Compliant
 * Operands of << operator (precedence 10) are:
 *   a precedence 16 does not need parentheses
 *   ( E ) precedence 16 already parenthesized
 */
x = a << ( b + c );

/* Compliant
 * Operands of && operator (precedence 4) are:
 *   a precedence 16 does not need parentheses
 *   && precedence 4 does not need parentheses
 */
if ( a && b && c )
{
}

/* Compliant
 * Operands of && operator (precedence 4) are:
 *   defined(X) precedence 14 does not need parentheses
 *   (E) precedence 16 already parenthesized
 */
#if defined ( X ) && ( ( X + Y ) > Z )

/* Compliant
 * Operands of && operator (precedence 4) are:
 *   !defined ( X ) precedence 14 does not need parentheses
 *   defined ( Y ) precedence 14 does not need parentheses
 *   Operand of ! operator (precedence 14) is:
 *   defined ( X ) precedence 14 does not need parentheses
 */
#if !defined ( X ) && defined ( Y )

```

Note: this rule does not require the operands of a `,` operator to be parenthesized. Use of the `,` operator is prohibited by Rule 12.3.

```
x = a, b; /* Compliant - parsed as ( x = a ), b */
```

Rule 12.2 The right hand operand of a shift operator shall lie in the range zero to one less than the width in bits of the *essential type* of the left hand operand

C90 [Undefined 32], C99 [Undefined 48], C11 [Undefined 51]

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Rationale

If the right hand operand is negative, or greater than or equal to the width of the left hand operand, then the behaviour is undefined.

If, for example, the left hand operand of a left-shift or right-shift is a 16-bit integer, then it is important to ensure that this is shifted only by a number in the range 0 to 15.

See Section 8.10 for a description of *essential type* and the limitations on the *essential types* for the operands of shift operators.

There are various ways of ensuring this rule is followed. The simplest is for the right hand operand to be a constant (whose value can then be statically checked). Use of an unsigned integer type will ensure that the operand is non-negative, so then only the upper limit needs to be checked (dynamically at run time or by review). Otherwise both limits will need to be checked.

Example

```
u8a = u8a << 7;          /* Compliant */
u8a = u8a << 8;          /* Non-compliant */
u16a = ( uint16_t ) u8a << 9; /* Compliant */
```

To assist in understanding the following examples, it should be noted that the *essential type* of `1u` is *essentially unsigned char*, whereas the *essential type* of `1UL` is *essentially unsigned long*.

```
1u << 10u;                /* Non-compliant */
( uint16_t ) 1u << 10u;   /* Compliant */
1UL << 10u;               /* Compliant */
```

Rule 12.3 The comma operator should not be used

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

Use of the comma operator is generally detrimental to the readability of code, and the same effect can usually be achieved by other means.

Example

```
f ( ( 1, 2 ), 3 ); /* Non-compliant - how many parameters? */
```

The following example is non-compliant with this rule and other rules:

```
for ( i = 0, p = &a[ 0 ]; i < N; ++i, ++p )
{
}
```

Rule 12.4 Evaluation of *constant expressions* should not lead to unsigned integer wrap-around

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

This rule applies to expressions that satisfy the *constraints* and semantics for a *constant expression*, whether or not they appear in a context that requires a *constant expression*.

If an expression is not evaluated, for example because it appears in the right operand of a logical AND operator whose left operand is always false, then this rule does not apply.

Rationale

Unsigned integer expressions do not strictly overflow, but instead wrap-around. Although there may be good reasons to use modulo arithmetic at run-time, it is less likely for its use to be intentional at compile-time.

Example

The expression associated with a *case* label is required to be a *constant expression*. If an unsigned wrap-around occurs during evaluation of a *case* expression, it is likely to be unintentional. On a machine with a 16-bit *int* type, any value of `BASE` greater than or equal to 65 024 would result in wrap-around in the following example:

```
#define BASE 65024u

switch ( x )
{
  case BASE + 0u:
    f ( );
    break;
  case BASE + 1u:
    g ( );
    break;
  case BASE + 512u: /* Non-compliant - wraps to 0 */
    h ( );
    break;
}
```

The controlling expression of a *#if* or *#elif* preprocessor directive is required to be a *constant expression*.

```
#if 1u + ( 0u - 10u ) /* Non-compliant as ( 0u - 10u ) wraps */
```

In this example, the expression `DELAY + WIDTH` has the value 70 000 but this will wrap-around to 4 464 on a machine with a 16-bit *int* type.

```
#define DELAY 10000u
#define WIDTH 60000u

void fixed_pulse ( void )
{
  uint16_t off_time16 = DELAY + WIDTH; /* Non-compliant */
}
```

This rule does not apply to the expression `c + 1` in the following compliant example as it accesses an object and therefore does not satisfy the semantics for a *constant expression*:

```
const uint16_t c = 0xffffu;

void f ( void )
{
  uint16_t y = c + 1u; /* Compliant */
}
```

In the following example, the sub-expression `(0u - 1u)` leads to unsigned integer wrap-around. In the initialization of `x`, the sub-expression is not evaluated and the expression is therefore compliant. However, in the initialization of `y`, it may be evaluated and the expression is therefore non-compliant.

```
bool_t b;

void g ( void )
{
  uint16_t x = ( 0u == 0u ) ? 0u : ( 0u - 1u ); /* Compliant */
  uint16_t y = b ? 0u : ( 0u - 1u ); /* Non-compliant */
}
```

Rule 12.5 The *sizeof* operator shall not have an operand which is a function parameter declared as “array of type”

Category	Mandatory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The function parameter `A` in `void f (int32_t A[4])` is declared as “array of type”.

Rationale

The *sizeof* operator can be used to determine the number of elements in an array `A`:

```
size_t arraySize = sizeof ( A ) / sizeof ( A[ 0 ] );
```

This works as expected when `A` is an identifier that designates an array, as it has type “array of type”. It does not “degenerate” to a pointer and `sizeof (A)` returns the size of the array.

However, this is not the case when `A` is a function parameter. The C Standard states that a function parameter never has type “array of type” and a function parameter declared as an array will “degenerate” to “pointer to type”. This means that `sizeof (A)` is equivalent to `sizeof (int32_t *)`, which does not return the size of an array.

Example

```
int32_t glbA[] = { 1, 2, 3, 4, 5 };

void f ( int32_t A[ 4 ] )
{
    /*
     * The following is non-compliant as it always gives the same answer,
     * irrespective of the number of members that appear to be in the array
     * (4 in this case), because A has type int32_t * and not int32_t[ 4 ].
     * As sizeof ( int32_t * ) is often the same as sizeof ( int32_t ),
     * numElements is likely to always have the value 1.
     */
    uint32_t numElements = sizeof ( A ) / sizeof ( int32_t );

    /*
     * The following is compliant as numElements_glbA will be given the
     * expected value of 5.
     */
    uint32_t numElements_glbA = sizeof ( glbA ) / sizeof ( glbA[ 0 ] );
}
```

Rule 12.6 Structure and union members of atomic objects shall not be directly accessed

C11 [Undefined 42]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C11

Amplification

The C Standard defines the following access functions for atomic objects: *atomic_init()*, *atomic_store()*, *atomic_load()*, *atomic_exchange()*, *atomic_compare_exchange()*.

Accesses to atomic objects of structure and union types shall only be made to the object as a whole, and only using these functions and the assignment operator =. In particular, the . and -> operators shall not be used on atomic objects of structure and union type.

Rationale

The Standard guarantees absence of data races when performing atomic operations on data shared between threads without requiring explicit protection via mutex or condition variables. The operations have to be performed by dedicated access functions which provide an appropriate built-in protection. Direct access to structure or union members of atomic objects circumvents this protection, thus making them vulnerable to data races.

Note: The *atomic_init()* functions does not avoid data races. Concurrent access to the variable being initialized, even via an atomic operation, constitutes a data race.

Example

```
typedef struct s {
    uint8_t a;
    uint8_t b;
} s_t;
_Atomic s_t astr;

sint32_t main(void)
{
    s_t lstr = { 7U, 42U };

    astr.b = 43U;                /* Non-compliant */

    lstr = atomic_load( &astr );
    lstr.b = 43U;
    atomic_store( &astr, lstr ); /* Compliant */

    lstr.a = 8U;
    astr = lstr;                /* Compliant */
}
```

See also

Dir 5.1, Rule 11.4, Rule 9.7

8.13 Side effects

Rule 13.1 *Initializer lists shall not contain persistent side effects*

C99 [Unspecified 17, 22], C11 [Unspecified 18, 23]

Category	Required
Analysis	Undecidable, System
Applies to	C99, C11

Rationale

C90 *constrains* the initializers for automatic objects with aggregate types to contain only *constant expressions*. However, later editions of the C Standard permit automatic aggregate initializers to contain expressions that are evaluated at run-time. It also permits compound literals which behave as anonymous initialized objects. The order in which *side effects* occur during evaluation of the expressions in an *initializer list* is unspecified and the behaviour of the initialization is therefore unpredictable if those *side effects* are *persistent*.

Example

```
volatile uint16_t v1;

void f ( void )
{
    /* Non-compliant - volatile access is persistent side effect */
    uint16_t a[ 2 ] = { v1, 0 };
}

void g ( uint16_t x, uint16_t y )
{
    /* Compliant - no side effects */
    uint16_t a[ 2 ] = { x + y, x - y };
}

uint16_t x = 0u;

extern void p ( uint16_t a[ 2 ] );

void h ( void )
{
    /* Non-compliant - two side effects */
    p ( ( uint16_t[ 2 ] ) { x++, x++ } );
}
```

See also

Rule 13.2

Rule 13.2 The value of an expression and its *persistent side effects* shall be the same under all permitted evaluation orders and shall be independent from thread interleaving

C90 [Unspecified 7–9; Undefined 18]
 C99 [Unspecified 15–18; Undefined 32]
 C11 [Unspecified 16–19; Undefined 35]

Category Required
Analysis Undecidable, System
Applies to C90, C99, C11

Amplification

Between any two adjacent sequence points:

1. No object shall be modified more than once;
2. All parts of the expression are considered when determining whether an object is read or modified, irrespective of any known values.
3. No object shall be both modified and read unless any such read of the object's value contributes towards computing the value to be stored into the object;
4. There shall be no more than one modification access with *volatile*-qualified or atomic type;
5. There shall be no more than one read access with *volatile*-qualified type;
6. There shall be no more than one read access to an object with atomic type.

Notes:

1. An object might be accessed indirectly, by means of a pointer or a called function, as well as being accessed directly by the expression.
2. This Amplification is intentionally stricter than the headline of the rule. As a result, expressions such as $x = x = 0;$ are not permitted by this rule even though the value and the *persistent side effects*, provided that x is not *volatile*, are independent of the order of evaluation or *side effects*.

Sequence points are summarized in Annex C of the C Standard. The sequence points in C90 are a subset of those in later editions.

Rationale

The C Standard gives considerable flexibility to compilers when evaluating expressions. Most operators can have their operands evaluated in any order. The main exceptions are:

- Logical AND `&&` in which the second operand is evaluated only if the first operand evaluates to non-zero;
- Logical OR `||` in which the second operand is evaluated only if the first operand evaluates to zero;

- The conditional operator `?:` in which the first operand is always evaluated and then either the second or third operand is evaluated;
- The `,` operator in which the first operand is evaluated and then the second operand is evaluated.

Note: The presence of parentheses may alter the order in which operators are applied. However, this does not affect the order of evaluation of the lowest-level operands, which may be evaluated in any order.

The atomic types provide assurance that a single read or write access to an atomic object is not subject to interruption or potential interference from other threads. However, that does not prevent two distinct atomic accesses to the same variable by a thread being pre-empted by another thread modifying that variable. On non-atomic variables such interference can only be caused by data races and constitute undefined behaviour. By definition, although there are no data races on atomic variables, such interference is still undesirable.

Many of the common instances of the unpredictable behaviour, associated with expression evaluation, can be avoided by following the advice given by this rule, by Rule 13.3, and by Rule 13.4. However, in order to simplify this rule, it does restrict some forms which are well-defined.

Examples

When the `COPY_ELEMENT` macro is invoked in this non-compliant example, `i` is read twice and modified twice. It is unspecified whether the order of operations on `i` is:

- Read, modify, read, modify, or
- Read, read, modify, modify.

```
#define COPY_ELEMENT( index ) ( a[( index )] = b[( index )] )
COPY_ELEMENT ( i++ );
```

In this non-compliant example the order in which `v1` and `v2` are read is unspecified.

```
extern volatile uint16_t v1, v2;
uint16_t t;

t = v1 + v2;
```

In this compliant example `PORT` is read and modified.

```
extern volatile uint8_t PORT;

PORT = PORT & 0x80u;
```

The order of evaluation of function arguments is unspecified as is the order in which *side effects* occur as shown in this non-compliant example.

```
uint16_t i = 0;

/*
 * Unspecified whether this call is equivalent to:
 *     f ( 0, 0 )
 * or   f ( 0, 1 )
 */
f ( i++, i );
```

The relative order of evaluation of a function designator and function arguments is unspecified. In this non-compliant example, if the call to `g` modifies `p` then it is unspecified whether the function designator `p->f` uses the value of `p` prior to the call of `g` or after it.

```
p->f ( g ( &p ) );
```

In the following example, Thread `T2` might interrupt Thread `T1` while the expression `a - a` is evaluated. Then the first load instruction for `a` loads the value 10, but the second load operation loads the value 7. The compliant solution avoids the problem by storing the value of `a` in a local variable.

```
_Atomic int32_t a;

int32_t t1( void* ignore ) /* Thread T1 entry */
{
    int32_t v1, v2;
    int32_t acopy;

    a      = 10;
    acopy = a; /* acopy may be either 10 or 7 */

    v1 = a - a; /* Non-compliant - v1 may be 0 or 3 */
    v2 = acopy - acopy; /* Compliant - v2 is always 0 */

    return v1 + v2;
}

int32_t t2( void* ignore ) /* Thread T2 entry */
{
    a = 7;

    return a;
}
```

See also

Dir 4.9, Rule 13.1, Rule 13.3, Rule 13.4

Rule 13.3 A *full expression* containing an increment (`++`) or decrement (`--`) operator should have no other potential *side effects* other than that caused by the increment or decrement operator

C90 [Unspecified 7, 8; Undefined 18]
 C99 [Unspecified 15; Undefined 32]
 C11 [Unspecified 16; Undefined 35]

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

A function call is considered to be a *side effect* for the purposes of this rule.

All sub-expressions of the *full expression* are treated as if they were evaluated for the purposes of this rule, even if specified as not being evaluated by the C Standard.

Rationale

The use of increment and decrement operators in combination with other operators is not recommended because:

- It can significantly impair the readability of the code;
- It introduces additional *side effects* into a statement with the potential for undefined behaviour (covered by Rule 13.2).

It is clearer to use these operations in isolation from any other operators.

Example

The expression:

```
u8a = u8b++
```

is non-compliant. The non-compliant expression statement:

```
u8a = ++u8b + u8c--;
```

is clearer when written as the following sequence:

```
++u8b;
u8a = u8b + u8c;
u8c--;
```

The following are all compliant because the only *side effect* in each expression is caused by the increment or decrement operator.

```
x++;
a[ i ]++;
b.x++;
c->x++;
++( *p );
*p++;
( *p )++;
```

The following are all non-compliant because they contain a function call as well as an increment or decrement operator:

```
if ( ( f ( ) + --u8a ) == 0u )
{
}

g ( u8b++ );
```

The following are all non-compliant even though the sub-expression containing the increment or decrement operator or some other *side effect* is not evaluated:

```
u8a = ( 1u == 1u ) ? 0u : u8b++;

if ( u8a++ == ( ( 1u == 1u ) ? 0u : f ( ) ) )
{
}
```

See also

Rule 13.2

Rule 13.4 The result of an assignment operator should not be *used*

C90 [Unspecified 7, 8; Undefined 18]
 C99 [Unspecified 15, 18; Undefined 32]
 C11 [Unspecified 16, 19; Undefined 35]
 [Koenig 6]

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule applies even if the expression containing the assignment operator is not evaluated.

Rationale

The use of assignment operators, simple or compound, in combination with other arithmetic operators is not recommended because:

- It can significantly impair the readability of the code;
- It introduces additional *side effects* into a statement making it more difficult to avoid the undefined behaviour covered by Rule 13.2.

Example

```
x = y;                /* Compliant */
a[ x ] = a[ x = y ]; /* Non-compliant - the value of x = y
                    * is used */

/*
 * Non-compliant - value of bool_var = false is used but
 * bool_var == false was probably intended
 */
if ( bool_var = false )
{
}

/* Non-compliant even though bool_var = true isn't evaluated */
if ( ( 0u == 0u ) || ( bool_var = true ) )
{
}

/* Non-compliant - value of x = f() is used */
if ( ( x = f ( ) ) != 0 )
{
}

/* Non-compliant - value of b += c is used */
a[ b += c ] = a[ b ];

/* Non-compliant - values of c = 0 and b = c = 0 are used */
a = b = c = 0;
```

See also

Rule 13.2

Rule 13.5 The right hand operand of a logical && or || operator shall not contain *persistent side effects*

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Rationale

The evaluation of the right-hand operand of the && and || operators is conditional on the value of the left-hand operand. If the right-hand operand contains *side effects* then those *side effects* may or may not occur which may be contrary to programmer expectations.

If evaluation of the right-hand operand would produce *side effects* which are not *persistent* at the point in the program where the expression occurs then it does not matter whether the right-hand operand is evaluated or not.

The term *persistent side effect* is defined in Appendix J.

Example

```

uint16_t f ( uint16_t y )
{
    /* These side effects are not persistent as seen by the caller */
    uint16_t temp = y;

    temp = y + 0x8080U;

    return temp;
}

uint16_t h ( uint16_t y )
{
    static uint16_t temp = 0;

    /* This side effect is persistent */
    temp = y + temp;

    return temp;
}

void g ( void )
{
    /* Compliant - f ( ) has no persistent side effects */
    if ( ishigh && ( a == f ( x ) ) )
    {
    }

    /* Non-compliant - h ( ) has a persistent side effect */
    if ( ishigh && ( a == h ( x ) ) )
    {
    }
}

volatile uint16_t v;
uint16_t x;

/* Non-compliant - access to volatile v is persistent */
if ( ( x == 0u ) || ( v == 1u ) )
{
}

/* Non-compliant if fp points to a function with persistent side effects */
( fp != NULL ) && ( *fp ) ( 0 );

```

Rule 13.6 The operand of the *sizeof* operator shall not contain any expression which has potential *side effects*

C99 [Unspecified 21], C11 [Unspecified 22]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

Any expressions appearing in the operand of a *sizeof* operator are not normally evaluated. This rule mandates that the evaluation of any such expression shall not contain *side effects*, whether or not it is actually evaluated.

A function call is considered to be a *side effect* for the purposes of this rule.

Rationale

The operand of a *sizeof* operator may be either an expression or may specify a type. If the operand contains an expression, a possible programming error is to expect that expression to be evaluated when it is actually not evaluated in most circumstances.

The C90 standard states that expressions appearing in the operand are not evaluated at run-time.

In C99 and later, expressions appearing in the operand are usually not evaluated at run-time. However, if the operand contains a variable-length array type then the array size expression will be evaluated if necessary. If the result can be determined without evaluating the array size expression then it is unspecified whether it is evaluated or not.

Exception

An expression of the form `sizeof (v)`, where `v` is an *lvalue* with a *volatile* qualified type that is not a variable-length array, is permitted.

Example

```
volatile int32_t i;
        int32_t j;
        size_t s;

s = sizeof ( j );           /* Compliant          */
s = sizeof ( j++ );       /* Non-compliant    */
s = sizeof ( i );         /* Compliant - exception */
s = sizeof ( int32_t );   /* Compliant          */
```

In this example the final *sizeof* expression illustrates how it is possible for a variable-length array size expression to have no effect on the size of the type. The operand is the type “array of *n* pointers to function with parameter type array of `v int32_t`”. Because the operand has variable-length array type, it is evaluated. However, the size of the array of *n* function pointers is unaffected by the parameter list for those function pointer types. Therefore, the *volatile*-qualified object `v` may or may not be evaluated and its *side effects* may or may not occur.

```
volatile uint32_t v;

void f ( int32_t n )
{
    size_t s;

    s = sizeof ( int32_t[ n ] );           /* Compliant */
    s = sizeof ( int32_t[ n++ ] );       /* Non-compliant */
    s = sizeof ( void ( *[ n ] ) ( int32_t a[ v ] ) ); /* Non-compliant */
}

```

See also

Rule 18.8

8.14 Control statement expressions

The term *loop counter* is used by some of the rules in this section. A *loop counter* is defined as an object, array element or member of a structure or union which satisfies the following:

1. It has a scalar type;
2. Its value varies monotonically on each iteration of a given instance of a loop; and
3. It is involved in a decision to exit the loop.

Note: the second condition means that the value of the *loop counter* must change on each iteration of the loop and it must always change in the same direction for a given instance of the loop. It may, however, change in different directions on different instances, for example sometimes reading the elements of an array backwards and sometimes reading them forwards.

According to this definition, a loop need not have just one *loop counter*: it is possible for a loop to have no *loop counter* or to have more than one. See Rule 14.2 for further restrictions on *loop counters* in *for* loops.

The following code fragments show examples of loops and their corresponding *loop counters*.

In this loop, *i* is a *loop counter* because it has a scalar type, varies monotonically (increasing) and is involved in the loop termination condition.

```
for ( uint16_t i = 0; a[ i ] < 10; ++i )
{
}

```

The following loop has no *loop counter*. The object *count* has scalar type and varies monotonically but is not involved in the termination condition.

```
extern volatile bool_t b;
uint16_t count = 0;

while ( b )
{
    count = count + 1U;
}

```

In the following code, both `i` and `sum` are scalar and monotonically varying (decreasing and increasing respectively). However `sum` is not a *loop counter* because it is not involved in the decision to exit the loop.

```
uint16_t sum = 0;

for ( uint16_t i = 10U; i != 0U; --i )
{
    sum += i;
}
```

In the following loop, `p` is a *loop counter*. It is not involved in the loop control expression but it is involved in a decision to exit the loop by means of the *break* statement.

```
extern volatile bool_t b;

extern char *p;

do
{
    if ( *p == '\0' )
    {
        break;
    }

    ++p;
} while ( b );
```

The *loop counter* in the following example is `p->count`.

```
struct s
{
    uint16_t count;
    uint16_t a[ 10 ];
};

extern struct s *p;

for ( p->count = 0U; p->count < 10U; ++( p->count ) )
{
    p->a[ p->count ] = 0U;
}
```

Rule 14.1 *A loop counter shall not have essentially floating type*

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Rationale

When using a floating-point *loop counter*, accumulation of rounding errors may result in a mismatch between the expected and actual number of iterations. This can happen when a loop step that is not a power of the floating-point radix is rounded to a value that can be represented.

Even if a loop with a floating-point *loop counter* appears to behave correctly on one implementation, it may give a different number of iterations on another implementation.

Example

In the following non-compliant example, the value of `counter` is unlikely to be 1 000 at the end of the loop.

```
uint32_t counter = 0u;

for ( float32_t f = 0.0f; f < 1.0f; f += 0.001f )
{
    ++counter;
}
```

The following compliant example uses an integer *loop counter* to guarantee 1 000 iterations and uses it to generate `f` for use within the loop.

```
float32_t f;

for ( uint32_t counter = 0u; counter < 1000u; ++counter )
{
    f = ( float32_t ) counter * 0.001f;
}
```

The following *while* loop is non-compliant because `f` is being used as a *loop counter*.

```
float32_t f = 0.0f;

while ( f < 1.0f )
{
    f += 0.001f;
}
```

The following *while* loop is compliant because `f` is not being used as a *loop counter*.

```
float32_t f;
uint32_t u32a;

f = read_float32 ( );

do
{
    u32a = read_u32 ( );
    /* f does not change in the loop so cannot be a loop counter */
} while ( ( ( float32_t ) u32a - f ) > 10.0f );
```

See also

Rule 14.2

Rule 14.2 A *for* loop shall be well-formed

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

The three clauses of a *for* statement are the:

First clause which

- Shall be empty, or
- Shall be an expression whose only *persistent side effect* is to set the value of the *loop counter*, or
- Shall define and initialize the *loop counter* (C99 and later).

Second clause which

- Shall be an expression that has no *persistent side effects*, and
- Shall use the *loop counter* and optionally *loop control flags*, and
- Shall not use any other object that is modified in the *for* loop body.

Third clause which

- Shall be an expression whose only *persistent side effect* is to modify the value of the *loop counter*, and
- Shall not use objects that are modified in the *for* loop body.

There shall only be one *loop counter* in a *for* loop, which shall not be modified in the *for* loop body.

A *loop control flag* is defined as a single identifier denoting an object with *essentially Boolean type* that is used in the *second clause*.

The behaviour of a *for* loop body includes the behaviour of any functions called within that statement.

Rationale

The *for* statement provides a general-purpose looping facility. Using a restricted form of loop makes code easier to review and to analyse.

Exception

All three clauses may be empty, for example `for (; ;)`, so as to allow for infinite loops.

Example

In the following example, *i* is the *loop counter* and *flag* is a *loop control flag*.

```
bool_t flag = false;
for ( int16_t i = 0; ( i < 5 ) && !flag; i++ )
{
    if ( C )
    {
        flag = true;           /* Compliant - allows early termination
                               * of loop                          */
    }

    i = i + 3;                /* Non-compliant - altering the loop
                               * counter                          */
}
```

See also

Rule 14.1, Rule 14.3, Rule 14.4

Rule 14.3 Controlling expressions shall not be invariant

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

This rule applies to:

- Controlling expressions of *if*, *while*, *for*, *do ... while* and *switch* statements;
- The first operand of the *?:* operator.

Rationale

If a controlling expression has an invariant value, it is possible that there is a programming error. Any code that cannot be reached due to the presence of an invariant expression may be removed by the compiler. This might have the effect of removing defensive code, for instance, from the executable.

Exception

1. Invariants that are used to create infinite loops are permitted.
2. A *do ... while* loop with an *essentially Boolean* controlling expression that evaluates to *false* and satisfies the *constraints* and semantics for an *integer constant expression* is permitted.

Example

```
s8a = ( u16a < 0u ) ? 0 : 1;   /* Non-compliant - u16a always >= 0 */
if ( u16a <= 0xffffu )
{
    /* Non-compliant - always true */
}
```

```

if ( 2 > 3 )
{
    /* Non-compliant - always false */
}

for ( s8a = 0; s8a < 130; ++s8a )
{
    /* Non-compliant - always true */
}

if ( ( s8a < 10 ) && ( s8a > 20 ) )
{
    /* Non-compliant - always false */
}

if ( ( s8a < 10 ) || ( s8a > 5 ) )
{
    /* Non-compliant - always true */
}

while ( s8a > 10 )
{
    if ( s8a > 5 )
    {
        /* Non-compliant - s8a not volatile */
    }
}

while ( true )
{
    /* Compliant by exception 1 */
}

do
{
    /* Compliant by exception 2 */
} while ( 0u == 1u );

const uint8_t numcyl = 4u;

/*
 * Non-compliant - compiler is permitted to assume that numcyl always
 * has value 4
 */
if ( numcyl == 4u )
{
}

const volatile uint8_t numcyl_cal = 4u;

/*
 * Compliant - compiler assumes numcyl_cal may be changed by
 * an external method, e.g. automotive calibration tool, even
 * though the program cannot modify its value
 */
if ( numcyl_cal == 4u )
{
}

uint16_t n;    /* 10 <= n <= 100 */
uint16_t sum;

sum = 0;

for ( uint16_t i = ( n - 6u ); i < n; ++i )
{
    sum += i;
}

```

```

if ( ( sum % 2u ) == 0u )
{
    /*
     * Non-compliant - sum is the sum of 6 consecutive non-negative
     * integers so must be an odd number. The controlling expression
     * of the if statement will always be false.
     */
}

do
{
    /* Non-compliant - not covered by exception 2 */
} while ( (s8a < 10) && (s8a > 20) );

```

See also

Rule 2.1, Rule 14.2

Rule 14.4 The controlling expression of an *if* statement and the controlling expression of an *iteration-statement* shall have *essentially Boolean* type

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The controlling expression of a *for* statement is optional. The rule does not require the expression to be present but **does** require it to have *essentially Boolean type* if it is present.

Rationale

Strong typing requires the controlling expression of an *if* statement or *iteration-statement* to have *essentially Boolean* type.

Example

```

int32_t *p, *q;

while ( p )           /* Non-compliant - p is a pointer */
{
}

while ( q != NULL )  /* Compliant */
{
}

while ( true )       /* Compliant */
{
}

extern bool_t flag;

while ( flag )       /* Compliant */
{
}

```

```

int32_t i;

if ( i )                /* Non-compliant          */
{
}

if ( i != 0 )          /* Compliant          */
{
}

```

See also

Rule 14.2, Rule 20.8

8.15 Control flow

Rule 15.1 The *goto* statement should not be used

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

Unconstrained use of *goto* can lead to programs that are unstructured and extremely difficult to understand.

In some cases a total ban on *goto* requires the introduction of flags to ensure correct control flow, and it is possible that these flags may themselves be less transparent than the *goto* they replace. Therefore, if this rule is not followed, the restricted use of *goto* is allowed where that use follows the guidance in Rule 15.2 and Rule 15.3.

See also

Rule 9.1, Rule 15.2, Rule 15.3, Rule 15.4

Rule 15.2 The *goto* statement shall jump to a label declared later in the same function

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

Unconstrained use of *goto* can lead to programs that are unstructured and extremely difficult to understand.

Restricting the use of *goto* so that “back” jumps are prohibited ensures that iteration only occurs if the iteration statements provided by the language are used, helping to minimize visual code complexity.

Example

```
void f ( void )
{
    int32_t j = 0;

L1:
    ++j;

    if ( 10 == j )
    {
        goto L2;          /* Compliant      */
    }

    goto L1;             /* Non-compliant */

L2:
    ++j;
}
```

See also

Rule 15.1, Rule 15.3, Rule 15.4

Rule 15.3 Any label referenced by a *goto* statement shall be declared in the same block, or in any block enclosing the *goto* statement

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

For the purposes of this rule, a *switch-clause* that does not consist of a compound statement is treated as if it were a block.

Rationale

Unconstrained use of *goto* can lead to programs that are unstructured and extremely difficult to understand.

Preventing jumps between blocks, or into nested blocks, helps to minimize visual code complexity.

Note: C99 and later are more restrictive when variably modified types are used. An attempt to make a jump from outside the scope of an identifier with a variably modified type into such a scope results in a *constraint* violation.

Example

```

void f1 ( int32_t a )
{
    if ( a <= 0 )
    {
        goto L2;          /* Non-compliant */
    }

    goto L1;             /* Compliant   */

    if ( a == 0 )
    {
        goto L1;        /* Compliant   */
    }

    goto L2;            /* Non-compliant */
}

L1:
    if ( a > 0 )
    {
        L2:
            ;
    }
}

```

In the following example, the label `L1` is defined in a block which encloses the block containing the `goto` statement. However, the jump crosses from one *switch-clause* to another and, since *switch-clauses* are treated like blocks for the purposes of this rule, the `goto` statement is non-compliant.

```

switch ( x )
{
    case 0:
        if ( x == y )
        {
            goto L1;
        }
        break;
    case 1:
        y = x;
L1:
        ++x;
        break;
    default:
        break;
}

```

See also

Rule 9.1, Rule 15.1, Rule 15.2, Rule 15.4, Rule 16.1

Rule 15.4 There should be no more than one *break* or *goto* statement used to terminate any iteration statement

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

Restricting the number of exits from a loop helps to minimize visual code complexity. The use of one *break* or *goto* statement allows a single secondary exit path to be created when early loop termination is required.

Example

Both of the following nested loops are compliant as each has a single *break* used for early loop termination.

```
for ( x = 0; x < LIMIT; ++x )
{
    if ( ExitNow ( x ) )
    {
        break;
    }

    for ( y = 0; y < x; ++y )
    {
        if ( ExitNow ( LIMIT - y ) )
        {
            break;
        }
    }
}
```

The following loop is non-compliant as there are multiple *break* and *goto* statements used for early loop termination.

```
for ( x = 0; x < LIMIT; ++x )
{
    if ( BreakNow ( x ) )
    {
        break;
    }
    else if ( GotoNow ( x ) )
    {
        goto EXIT;
    }
    else
    {
        KeepGoing ( x );
    }
}

EXIT:
;
```

In the following example, the inner *while* loop is compliant because there is a single *goto* statement that can cause its early termination. However, the outer *while* loop is non-compliant because it can be terminated early either by the *break* statement or by the *goto* statement in the inner *while* loop.

```
while ( x != 0u )
{
    x = calc_new_x ( );

    if ( x == 1u )
    {
        break;
    }

    while ( y != 0u )
    {
        y = calc_new_y ( );

        if ( y == 1u )
        {
            goto L1;
        }
    }
}

L1:
z = x + y;
```

See also

Rule 15.1, Rule 15.2, Rule 15.3

Rule 15.5 A function should have a single point of exit at the end

[IEC 61508-3 Table B.9], [ISO 26262-6 Table 8]

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

A function should have no more than one *return* statement.

When a *return* statement is used, it should be the final statement in the compound statement that forms the body of the function.

Rationale

A single point of exit is required by IEC 61508 and ISO 26262 as part of the requirements for a modular approach.

Early returns may lead to the unintentional omission of function termination code.

If a function has exit points interspersed with statements that produce *persistent side effects*, it is not easy to determine which *side effects* will occur when the function is executed.

Example

In the following non-compliant code example, early returns are used to validate the function parameters.

```
bool_t f ( uint16_t n, char *p )
{
    if ( n > MAX )
    {
        return false;
    }

    if ( p == NULL )
    {
        return false;
    }

    return true;
}
```

See also

Rule 17.4

Rule 15.6 The body of an *iteration-statement* or a *selection-statement* shall be a *compound-statement*

[Koenig 24]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The body of an *iteration-statement* (*while*, *do ... while* or *for*) or a *selection-statement* (*if*, *else*, *switch*) shall be a *compound-statement*.

Rationale

It is possible for a developer to mistakenly believe that a sequence of statements forms the body of an *iteration-statement* or *selection-statement* by virtue of their indentation. The accidental inclusion of a semi-colon after the controlling expression is a particular danger, leading to a null control statement. Using a *compound-statement* clearly defines which statements actually form the body.

Additionally, it is possible that indentation may lead a developer to associate an *else* statement with the wrong *if*.

Exception

An *if* statement immediately following an *else* need not be contained within a *compound-statement*.

Example

The layout for the *compound-statement* and its enclosing braces are style issues which are not addressed by this document; the style used in the following examples is not mandatory.

Maintenance to the following

```
while ( data_available )
    process_data ( );          /* Non-compliant */
```

could accidentally give

```
while ( data_available )
    process_data ( );          /* Non-compliant */
    service_watchdog ( );
```

where `service_watchdog()` should have been added to the loop body. The use of a *compound-statement* significantly reduces the chance of this happening.

The next example appears to show that `action_2()` is the *else* statement to the first *if*.

```
if ( flag_1 )
    if ( flag_2 )                /* Non-compliant */
        action_1 ( );          /* Non-compliant */
else
    action_2 ( );                /* Non-compliant */
```

when the actual behaviour is

```
if ( flag_1 )
{
    if ( flag_2 )
    {
        action_1 ( );
    }
    else
    {
        action_2 ( );
    }
}
```

The use of *compound-statements* ensures that *if* and *else* associations are clearly defined.

The exception allows the use of *else if*, as shown below

```
if ( flag_1 )
{
    action_1 ( );
}
else if ( flag_2 )                /* Compliant by exception */
{
    action_2 ( );
}
else
{
    ; /* no action */
}
```

The following example shows how a spurious semi-colon could lead to an error

```
while ( flag );                  /* Non-compliant */
{
    flag = fn ( );
}
```

The following example shows the compliant method of writing a loop with an empty body:

```
while ( !data_available )
{
}
```

Rule 15.7 All *if ... else if* constructs shall be terminated with an *else* statement

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

A final *else* statement shall always be provided whenever an *if* statement is followed by a sequence of one or more *else if* constructs. The *else* statement shall contain at least either one *side effect* or a comment.

A function call is considered to be a *side effect* for the purposes of this rule.

Rationale

Terminating a sequence of *if ... else if* constructs with an *else* statement is defensive programming and complements the requirement for a *default* clause in a *switch* statement (see Rule 16.5).

The *else* statement is required to have a *side effect* or a comment to ensure that a positive indication is given of the desired behaviour, aiding the code review process.

Note: a final *else* statement is not required for a simple *if* statement.

Example

The following example is non-compliant as there is no explicit indication that no action is to be taken by the terminating *else*.

```
if ( flag_1 )
{
    action_1 ( );
}
else if ( flag_2 )
{
    action_2 ( );
}

/* Non-compliant */
```

The following shows a compliant terminating *else*.

```
else
{
    ; /* No action required - ; is optional */
}
```

See also

Rule 16.4

8.16 Switch statements

Rule 16.1 All *switch* statements shall be well-formed

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

A *switch* statement shall be considered to be well-formed if it conforms to the subset of C *switch* statements that is specified by the following syntax rules. If a syntax rule given here has the same name as one defined in the C Standard then it replaces the standard version for the scope of the *switch* statement; otherwise, all syntax rules given in the C Standard are unchanged.

switch-statement:

```
switch ( switch-expression ) { case-label-clause-list final-default-clause-list }
switch ( switch-expression ) { initial-default-clause-list case-label-clause-list }
```

case-label-clause-list:

```
case-clause-list
case-label-clause-list case-clause-list
```

case-clause-list:

```
case-label switch-clause
case-label case-clause-list
```

case-label:

```
case constant-expression:
```

final-default-clause-list:

```
default: switch-clause
case-label final-default-clause-list
```

initial-default-clause-list:

```
default: switch-clause
default: case-clause-list
```

switch-clause:

```
statement-listopt break;
C90: { declaration-listopt statement-listopt break; }
C99 and later: { block-item-listopt break; }
```

Except where explicitly permitted by this syntax, the *case* and *default* keywords may not appear anywhere within a *switch* statement body.

Note: some of the restrictions imposed on *switch* statements by this rule are expounded in the rules referenced in the "See also" section. It is therefore possible for code to violate both this rule **and** one of the more specific rules.

Note: the term ***switch label*** is used within the text of the specific *switch* statement rules to denote either a *case* label or a *default* label.

Rationale

The syntax for the *switch* statement in C is not particularly rigorous and can allow complex, unstructured behaviour. This and other rules impose a simple and consistent structure on the *switch* statement.

Example

The remaining rules in this section give examples that are also relevant to this rule.

See also

Rule 15.3, Rule 16.2, Rule 16.3, Rule 16.4, Rule 16.5, Rule 16.6

Rule 16.2 A *switch label* shall only be used when the most closely-enclosing compound statement is the body of a *switch* statement

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

The C Standard permits a *switch label*, i.e. a *case* label or *default* label, to be placed before any statement contained in the body of a *switch* statement, potentially leading to unstructured code. In order to prevent this, a *switch label* shall only appear at the outermost level of the compound statement forming the body of a *switch* statement.

Example

```
switch ( x )
{
  case 1:          /* Compliant      */
    if ( flag )
    {
  case 2:          /* Non-compliant */
    x = 1;
    }
    break;
  default:
    break;
}
```

See also

Rule 16.1

Rule 16.3 An unconditional *break* statement shall terminate every *switch-clause*

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

If a developer fails to end a *switch-clause* with a *break* statement, then control flow “falls” into the following *switch-clause* or, if there is no such clause, off the end and into the statement following the *switch* statement. Whilst falling into a following *switch-clause* is sometimes intentional, it is often an error. An unterminated *switch-clause* occurring at the end of a *switch* statement may fall into any *switch-clauses* which are added later.

To ensure that such errors can be detected, the last statement in every *switch-clause* shall be a *break* statement, or if the *switch-clause* is a compound statement, the last statement in the compound statement shall be a *break* statement.

Note: a *switch-clause* is defined as containing at least one statement. Two consecutive labels, *case* or *default*, do not have any intervening statement and are therefore permitted by this rule.

Example

```
switch ( x )
{
  case 0:
    break;           /* Compliant - unconditional break          */
  case 1:
    /* Compliant - empty fall through allows a group */
  case 2:
    break;           /* Compliant          */
  case 4:
    a = b;           /* Non-compliant - break omitted          */
  case 5:
    if ( a == b )
    {
      ++a;
      break;         /* Non-compliant - conditional break      */
    }
  default:
    ;                 /* Non-compliant - default must also have a break */
}
```

See also

Rule 16.1

Rule 16.4 Every *switch* statement shall have a *default* label

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The *switch-clause* following the *default* label shall, prior to the terminating *break* statement, contain either:

- A statement, or
- A comment.

Rationale

The requirement for a *default* label is defensive programming. Any statements following the *default* label are intended to take some appropriate action. If no statements follow the label then the comment can be used to explain why no specific action has been taken.

Example

```
int16_t x;

switch ( x )
{
  case 0:
    ++x;
    break;
  case 1:
  case 2:
    break;
}
/* Non-compliant - default label is required */

int16_t x;

switch ( x )
{
  case 0:
    ++x;
    break;
  case 1:
  case 2:
    break;
  default:
    /* Compliant - default label is present */
    errorflag = 1; /* should be non-empty if possible */
    break;
}
```

```

enum Colours
{ RED, GREEN, BLUE } colour;

switch ( colour )
{
  case RED:
    next = GREEN;
    break;
  case GREEN:
    next = BLUE;
    break;
  case BLUE:
    next = RED;
    break;

    /* Non-compliant - no default label.
    * Even though all values of the enumeration are
    * handled there is no guarantee that colour takes
    * one of those values */
}

```

See also

Rule 2.1, Rule 16.1

Rule 16.5 A *default* label shall appear as either the first or the last *switch label* of a *switch* statement

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

This rule makes it easy to locate the *default* label within a *switch* statement.

Example

```

switch ( x )
{
  default: /* Compliant - default is the first label */
  case 0:
    ++x;
    break;
  case 1:
  case 2:
    break;
}

switch ( x )
{
  case 0:
    ++x;
    break;
  default: /* Non-compliant - default is mixed with the case labels */
    x = 0;
    break;
  case 1:
  case 2:
    break;
}

```

```

switch ( x )
{
  case 0:
    ++x;
    break;
  case 1:
  case 2:
    break;
  default: /* Compliant - default is the final label */
    x = 0;
    break;
}

```

See also

Rule 15.7, Rule 16.1

Rule 16.6 Every *switch* statement shall have at least two *switch-clauses*

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

A *switch* statement with a single path is redundant and may be indicative of a programming error.

Example

```

switch ( x )
{
  default: /* Non-compliant - switch is redundant */
    x = 0;
    break;
}

switch ( y )
{
  case 1:
  default: /* Non-compliant - switch is redundant */
    y = 0;
    break;
}

switch ( z )
{
  case 1:
    z = 2;
    break;
  default: /* Compliant */
    z = 0;
    break;
}

```

See also

Rule 16.1

Rule 16.7 A *switch-expression* shall not have essentially Boolean type

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

The C Standard requires the controlling expression of a *switch* statement to have an integer type. Since the type that is used to implement Boolean values is an integer, it is possible to have a *switch* statement controlled by a Boolean expression. In this instance an *if-else* construct would be more appropriate.

Example

```
switch ( x == 0 ) /* Non-compliant - essentially Boolean */
{ /* In this case an "if-else" would be more logical */
  case false:
    y = x;
    break;
  default:
    y = z;
    break;
}
```

8.17 Functions

Rule 17.1 The standard *header file* `<stdarg.h>` shall not be used

C90 [Undefined 45, 70–76]
 C99 [Undefined 81, 128–135]
 C11 [Undefined 87, 136–143]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The standard *header file* `<stdarg.h>` shall not be `#include'd`, and none of the features that are specified as being provided by `<stdarg.h>` shall be used.

Rationale

The C Standard lists many instances of undefined behaviour associated with the features of `<stdarg.h>`, including:

- `va_end` not being used prior to end of a function in which `va_start` was used;
- `va_arg` being used in different functions on the same `va_list`;
- The type of an argument not being compatible with the type specified to `va_arg`.

Example

```
#include <stdarg.h>

void h ( va_list ap )          /* Non-compliant          */
{
    double y;

    y = va_arg ( ap, double ); /* Non-compliant          */
}

void f ( uint16_t n, ... )
{
    uint32_t x;

    va_list ap;                /* Non-compliant          */

    va_start ( ap, n );        /* Non-compliant          */
    x = va_arg ( ap, uint32_t ); /* Non-compliant          */

    h ( ap );

    /* undefined - ap is indeterminate because va_arg used in h ( ) */
    x = va_arg ( ap, uint32_t ); /* Non-compliant          */

    /* undefined - returns without using va_end ( ) */
}

void g (void )
{
    /* undefined - uint32_t:double type mismatch when f uses va_arg ( ) */
    f ( 1, 2.0, 3.0 );
}
```

Rule 17.2 Functions shall not call themselves, either directly or indirectly

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Rationale

Recursion carries with it the danger of exceeding available stack space, which can lead to failure. Unless recursion is very tightly controlled, it is not possible to determine before execution what the worst-case stack usage could be.

Rule 17.3 A function shall not be declared implicitly

C90 [Undefined 6, 22, 23]

Category	Mandatory
Analysis	Decidable, Single Translation Unit
Applies to	C90

Rationale

Provided that a function call is made in the presence of a prototype, a *constraint* ensures that the number of arguments matches the number of parameters and that each argument can be assigned to its corresponding parameter.

If a function is declared implicitly, a C90 compiler will assume that the function has a return type of *int*. Since an implicit function declaration does not provide a prototype, a compiler will have no information about the number of function parameters and their types. Inappropriate type conversions may result in passing the arguments and assigning the return value, as well as other undefined behaviour.

Example

If the function `power` is declared as:

```
extern double power ( double d, int n );
```

but the declaration is **not** visible in the following code then undefined behaviour will occur.

```
void func ( void )
{
    /* Non-compliant - return type and both argument types incorrect */
    double sq1 = power ( 1, 2.0 );
}
```

See also

Rule 8.2, Rule 8.4

Rule 17.4 All exit paths from a function with non-*void* return type shall have an explicit *return* statement with an expression

C90 [Undefined 43], C99 [Undefined 82], C11 [Undefined 88]

Category Mandatory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

The expression given to the *return* statement provides the value that the function returns. If a non-*void* function does not return a value but the calling function uses the returned value, the behaviour is undefined. This can be avoided by ensuring that, in a non-*void* function:

- Every *return* statement has an expression, and
- Control cannot reach the end of the function without encountering a *return* statement.

Note: C99 and later *constrain* every *return* statement in a non-*void* function to return a value.

Exception

For C99 and later, the C Standard specifies that if control reaches the end of `main` without encountering a `return` statement, the effect is that of executing `return 0`. Therefore, for C99 and later, the `return` statement may be omitted for function `main`.

Example

```
int32_t absolute ( int32_t v )
{
    if ( v < 0 )
    {
        return v;
    }

    /*
     * Non-compliant - control can reach this point without
     * returning a value
     */
}

uint16_t lookup ( uint16_t v )
{
    if ( ( v < V_MIN ) || ( v > V_MAX ) )
    {
        /* Non-compliant - no value returned. Constraint in C99 and later */
        return;
    }

    return table[ v ];
}
```

See also

Rule 15.5

Rule 17.5 The function argument corresponding to a parameter declared to have an array type shall have an appropriate number of elements

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

If a parameter is declared as an array with a specified size, the corresponding argument in each function call shall point into an object that has at least as many elements as the array.

Rationale

The use of an array declarator for a function parameter specifies the function interface more clearly than using a pointer. The minimum number of elements expected by the function is explicitly stated, whereas this is not possible with a pointer.

A function parameter array declarator which does not specify a size is assumed to indicate that the function can handle an array of any size. In such cases, it is expected that the array size will be communicated by some other means, for example by being passed as another parameter, or by terminating the array with a sentinel value.

The use of an array bound is recommended as it allows out-of-bounds checking to be implemented within the function body and extra checks on parameter passing. It is legal in C to pass an array of the incorrect size to a parameter with a specified size, which can lead to unexpected behaviour.

Example

```

/*
 * Intent is that function does not access outside the range
 * array1[ 0 ] .. array1[ 3 ]
 */
void fn1 ( int32_t array1[ 4 ] );

/* Intent is that function handles arrays of any size */
void fn2 ( int32_t array2[ ] );

void fn ( int32_t *ptr )
{
    int32_t arr3[ 3 ] = { 1, 2, 3 };
    int32_t arr4[ 4 ] = { 0, 1, 2, 3 };

    /* Compliant - size of array matches the prototype */
    fn1 ( arr4 );

    /* Non-compliant - size of array does not match prototype */
    fn1 ( arr3 );

    /* Compliant only if ptr points to at least 4 elements */
    fn1 ( ptr );

    /* Compliant */
    fn2 ( arr4 );

    /* Compliant */
    fn2 ( ptr );
}

```

See also

Rule 17.6

Rule 17.6 The declaration of an array parameter shall not contain the *static* keyword between the []

C99 [Undefined 71], C11 [Undefined 77]

Category Mandatory

Analysis Decidable, Single Translation Unit

Applies to C99, C11

Rationale

The C Standard provides a mechanism for the programmer to inform the compiler that an array parameter contains a specified minimum number of elements. Some compilers are able to take advantage of this information to generate more efficient code for some types of processor.

If the guarantee made by the programmer is not honoured, and the number of elements is less than the minimum specified, the behaviour is undefined.

The processors used in typical embedded applications are unlikely to provide the facilities required to take advantage of the additional information provided by the programmer. The risk of the program failing to meet the guaranteed minimum number of elements outweighs any potential performance increase.

Example

There is no use of this language feature that is compliant with this rule. The examples show some of the undefined behaviour that can arise from its use.

```

/* Non-compliant - uses static in array declarator */
uint16_t total ( uint16_t n, uint16_t a[ static 20 ] )
{
    uint16_t i;

    uint16_t sum = 0U;

    /* Undefined behaviour if a has fewer than 20 elements */
    for ( i = 0U; i < n; ++i )
    {
        sum = sum + a[ i ];
    }

    return sum;
}

extern uint16_t v1[ 10 ];
extern uint16_t v2[ 20 ];

void g ( void )
{
    uint16_t x;

    x = total ( 10U, v1 );    /* Undefined - v1 has 10 elements but needs
                             *                               at least 20 */
    x = total ( 20U, v2 );  /* Defined but non-compliant */
}

```

See also

Rule 17.5

Rule 17.7 The value returned by a function having non-*void* return type shall be *used*

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

It is possible to call a function without using the return value, which may be an error. If the return value of a function is intended not to be *used* explicitly, it should be cast to the *void* type. This has the effect of using the value without violating Rule 2.2.

Example

```

uint16_t func ( uint16_t para1 )
{
    return para1;
}

```

```

uint16_t x;

void discarded ( uint16_t para2 )
{
    func ( para2 );           /* Non-compliant - value discarded */
    ( void ) func ( para2 ); /* Compliant */
    x = func ( para2 );      /* Compliant */
}

```

See also

Dir 4.7, Rule 2.2

Rule 17.8 A function parameter should not be modified

Category Advisory

Analysis Undecidable, System

Applies to C90, C99, C11

Rationale

A function parameter behaves in the same manner as an object that has automatic storage duration. While the C language permits parameters to be modified, such use can be confusing and conflict with programmer expectations. It may be less confusing to copy the parameter to an automatic object and modify that copy. With a modern compiler, this will not usually result in any storage or execution time penalty.

Programmers who are unfamiliar with C, but who are used to other languages, may modify a parameter believing that the effects of the modification will be felt in the calling function.

Example

```

int16_t glob = 0;

void proc ( int16_t para )
{
    para = glob;           /* Non-compliant */
}

void f ( char *p, char *q )
{
    p = q;                 /* Non-compliant */
    *p = *q;               /* Compliant */
}

```

Rule 17.9 A function declared with a `_Noreturn` function specifier shall not return to its caller

C11 [Undefined 71]

Category	Mandatory
Analysis	Undecidable, System
Applies to	C11

Rationale

By definition, use of the `_Noreturn` function specifier indicates unambiguously that a function is not expected to return to its caller, by any path.

Returning from such a function is indicative of an error in the program's control flow, resulting in undefined behaviour.

Example

```
_Noreturn void a ( void )
{
    return; /* Non-compliant - breaks _Noreturn contract */
}

_Noreturn void b ( void )
{
    while ( true ) /* Permitted by Rule 14.3 Exception 1 */
    {
        ...
    }

    /* Compliant - function can never return */
}

_Noreturn void c ( void )
{
    abort();          /* Compliant - no return from abort() */
}

_Noreturn void d ( int32_t i )
{
    if ( i > 0 )
    {
        abort();
    }

    /* Non-compliant - returns if i <= 0 */
}
```

See also

Rule 17.11

Rule 17.10 A function declared with a `_Noreturn` function specifier shall have `void` return type

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C11

Rationale

By definition, use of the `_Noreturn` function specifier indicates unambiguously that a function will not return to its caller.

A function that cannot return cannot provide a return value and shall therefore be defined with a `void` return type.

Example

```
_Noreturn int32_t f ( void ); /* Non-compliant */
_Noreturn void g ( void ); /* Compliant */
```

See also

Rule 17.9, Rule 17.11

Rule 17.11 A function that never returns should be declared with a `_Noreturn` function specifier

Category	Advisory
Analysis	Undecidable, System
Applies to	C11

Rationale

Declaring a function that cannot return as `_Noreturn` highlights that this is “by design”.

Exception

This rule does not apply to `main()`, as the C Standard states that it is a constraint violation for `main()` to be declared with the `_Noreturn` function specifier.

Example

```
void f ( void ) /* Non-compliant - exit() call means there is no return */
{
    exit ( 0 );
}
```

See also

Rule 17.9

Rule 17.12 A function identifier should only be used with either a preceding `&`, or with a parenthesized parameter list

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

A function identifier not preceded by `&` and not followed by a parenthesized parameter list (which may be empty) can be confusing: it may not be clear whether the intent is to call the function (but the parenthesized parameter list has accidentally been omitted) or the intent is to obtain the function address.

Example

```
typedef int32_t (*pfn_i)(void);

extern int32_t func1 ( void ); /* Note: A function */
extern int32_t (*func2)( void ); /* Note: A function pointer */

void func ( void )
{
    pfn_i pfn1 = &func1; /* Compliant */
    pfn_i pfn2 = func1; /* Non-compliant */

    int32_t i32a = (*pfn1)(); /* Compliant - explicit call via a fn-pointer */
    pfn1(); /* Compliant - implicit call via a fn-pointer */

    if ( func1 == func2 ) /* Non-compliant */
    {
        /* ... */
    }

    if ( func1 ( ) == func2 ( ) ) /* Compliant - comparing return values */
    {
        /* ... */
    }

    if ( &func1 == func2 ) /* Compliant - comparing a function's address */
    {
        /* ... */
    }
}
```

Rule 17.13 A *function type* shall not be type qualified

C90 [Undefined *], C99 [Undefined 63], C11 [Undefined 66]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The type qualifiers are *const*, *volatile*, *restrict* or *_Atomic*.

Rationale

The behaviour is undefined if the specification of a *function type* includes any type qualifiers.

Note: this rule applies to a *function type* and not to the return type of a function.

Example

```
const uint16_t    cf (void);          /* Compliant    - returns const uint16_t    */
const uint16_t * pcf (void);        /* Compliant    - returns a pointer to
                                     const uint16_t */
```

In the following examples, `ftype` is the type of a function returning `uint16t`:

```
typedef uint16_t ftype (void);

typedef const ftype      cftype; /* Non-compliant - cftype is const-qualified */

typedef      ftype      dftype; /* Compliant    - dftype is not qualified    */
typedef      ftype * const pcftype; /* Compliant    - const pointer to ftype    */

typedef const uint16_t * cfptype (void); /* Compliant - cfptype is the type of a
                                     function returning
                                     const uint16_t * */
```

8.18 Pointers and arrays

Rule 18.1 A pointer resulting from arithmetic on a pointer operand shall address an element of the same array as that pointer operand

C90 [Undefined 30], C99 [Undefined 43, 44, 46, 59], C11 [Undefined 46, 47, 49, 62]

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

Creation of a pointer to one beyond the end of the array is well-defined by the C Standard and is permitted by this rule. Dereferencing a pointer to one beyond the end of an array results in undefined behaviour and is forbidden by this rule.

This rule applies to all forms of array indexing:

```
integer_expression + pointer_expression
pointer_expression + integer_expression
pointer_expression - integer_expression
pointer_expression += integer_expression
pointer_expression -= integer_expression
++pointer_expression
pointer_expression++
--pointer_expression
pointer_expression--
pointer_expression [ integer_expression ]
integer_expression [ pointer_expression ]
```

Notes:

1. A subarray is also an array.
2. Within the description of the semantics for the “additive operators”, the C Standard states that, for the purposes of pointer arithmetic, a pointer to an object that is not an array is treated as if it were a pointer to the first element of an array with a single element.
3. A pointer to an object of type T which has been converted to a pointer to an object of type `char`, `signed char` or `unsigned char` (see exception to Rule 11.3) is treated as an array of that type with bound equal to `sizeof(T)`.

Rationale

Although some compilers may be able to determine at compile time that an array boundary has been exceeded, no checks are generally made at run-time for invalid array subscripts. Using an invalid array subscript can lead to erroneous behaviour of the program.

Run-time derived array subscript values are of the most concern since they cannot easily be checked by static analysis or manual review. Code of a defensive programming nature should, where possible and practicable, be provided to check such subscript values against valid ones and, if required, appropriate action be taken.

It is undefined behaviour if the result obtained from one of the above expressions is not a pointer to an element of the array pointed to by `pointer_expression` or an element one beyond the end of that array.

Multi-dimensional arrays are “arrays of arrays”. This rule does not allow pointer arithmetic that results in the pointer addressing a different subarray. Array subscripting over “internal” boundaries shall not be used, as such behaviour is undefined.

Example

The use of the `+` operator will also violate Rule 18.4.

```
int32_t f1 ( int32_t * const a1, int32_t a2[ 10 ] )
{
    int32_t *p = &a1[ 3 ];           /* Compliant/non-compliant depending on
                                     * the value of a1 */
    return *( a2 + 9 );             /* Compliant */
}

void f2 ( void )
{
    int32_t data = 0;
    int32_t b = 0;
    int32_t c[ 10 ] = { 0 };
    int32_t d[ 5 ][ 2 ] = { 0 }; /* 5-element array of 2-element arrays
                                     * of int32_t */

    int32_t *p1 = &c[ 0 ];           /* Compliant */

    int32_t *p2 = &c[ 10 ];          /* Compliant - points to one beyond */

    int32_t *p3 = &c[ 11 ];          /* Non-compliant - undefined, points to
                                     * two beyond */
    data = *p2;                     /* Non-compliant - undefined, dereference
                                     * one beyond */

    data = f1 ( &b, c );
    data = f1 ( c, c );
}
```

```

p1++; /* Compliant */
c[ -1 ] = 0; /* Non-compliant - undefined, array
* bounds exceeded */

data = c[ 10 ]; /* Non-compliant - undefined, dereference
* of address one beyond */
data = *( &data + 0 ); /* Compliant - C treats data as an
* array of size 1 */

d[ 3 ][ 1 ] = 0; /* Compliant */
data = *( *( d + 3 ) + 1 ); /* Compliant */
data = d[ 2 ][ 3 ]; /* Non-compliant - undefined, internal
* boundary exceeded */

p1 = d[ 1 ]; /* Compliant */
data = p1[ 1 ]; /* Compliant - p1 addresses an array
* of size 2 */
}

```

The following example illustrates pointer arithmetic applied to members of a structure. Because each member is an object in its own right, this rule prevents the use of pointer arithmetic to move from one member to the next. However, it does not prevent arithmetic on a pointer to a member provided that the resulting pointer remains within the bounds of the member object.

```

struct
{
    uint16_t x;
    uint16_t y;
    uint16_t z;
    uint16_t a[ 10 ];
} s;

uint16_t *p;

void f3 ( void )
{
    p = &s.x;
    ++p; /* Compliant - p points one beyond s.x */
    p[ 0 ] = 1; /* Non-compliant - undefined, dereference of address one
* beyond s.x which is not necessarily
* the same as s.y */
    p[ 1 ] = 2; /* Non-compliant - undefined */

    p = &s.a[ 0 ]; /* Compliant - p points into s.a */
    p = p + 8; /* Compliant - p still points into s.a */
    p = p + 3; /* Non-compliant - undefined, p points more than one
* beyond s.a */
}

```

See also

Dir 4.1, Rule 11.3, Rule 18.4

Rule 18.2 Subtraction between pointers shall only be applied to pointers that address elements of the same array

C90 [Undefined 31], C99 [Undefined 45], C11 [Undefined 48]

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Rationale

This rule applies to expressions of the form

```
pointer_expression_1 - pointer_expression_2
```

It is undefined behaviour if `pointer_expression_1` and `pointer_expression_2` do not point to elements of the same array or the element one beyond the end of that array.

Example

```
#include <stddef.h>

void f1 ( int32_t *ptr )
{
    int32_t  a1[ 10 ];
    int32_t  a2[ 10 ];
    int32_t  *p1 = &a1[ 1 ];
    int32_t  *p2 = &a2[ 10 ];
    ptrdiff_t diff;

    diff = p1 - a1;          /* Compliant */
    diff = p2 - a2;          /* Compliant */
    diff = p1 - p2;          /* Non-compliant */
    diff = ptr - p1;          /* Non-compliant */
}
```

See also

Dir 4.1, Rule 18.4

Rule 18.3 The relational operators `>`, `>=`, `<` and `<=` shall not be applied to expressions of pointer type except where they point into the same object

C90 [Undefined 33], C99 [Undefined 50], C11 [Undefined 53]

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Rationale

Attempting to make comparisons between pointers will produce undefined behaviour if the two pointers do not point to the same object.

Note: it is permissible to address the next element beyond the end of an array, but accessing this element is not allowed.

Example

```
void f1 ( void )
{
    int32_t  a1[ 10 ];
    int32_t  a2[ 10 ];
    int32_t  *p1 = a1;

    if ( p1 < a1 )          /* Compliant      */
    {
    }
    if ( p1 < a2 )          /* Non-compliant */
    {
    }
}

struct limits
{
    int32_t lwb;
    int32_t upb;
};

void f2 ( void )
{
    struct limits limits_1 = { 2, 5 };
    struct limits limits_2 = { 10, 5 };

    if ( &limits_1.lwb <= &limits_1.upb ) /* Compliant      */
    {
    }
    if ( &limits_1.lwb > &limits_2.upb ) /* Non-Compliant */
    {
    }
}
```

See also

Dir 4.1

Rule 18.4 The `+`, `-`, `+=` and `-=` operators should not be applied to an expression of pointer type

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

Array indexing using the array subscript syntax, `ptr[expr]`, is the preferred form of pointer arithmetic because it is often clearer and hence less error prone than pointer manipulation. Any explicitly calculated pointer value has the potential to access unintended or invalid memory addresses. Such behaviour is also possible with array indexing, but the subscript syntax may ease the task of manual review.

Pointer arithmetic in C can be confusing to the novice. The expression `ptr+1` may be mistakenly interpreted as the addition of 1 to the address held in `ptr`. In fact the new memory address depends on the size in bytes of the pointer's target. This misunderstanding can lead to unexpected behaviour if `sizeof` is applied incorrectly.

When used with caution however, pointer manipulation using `++` can in some cases be considered more natural; e.g. sequentially accessing locations during a memory test where it is more convenient to treat the memory space as a contiguous set of locations and the address bounds can be determined at compilation time.

Exception

Subject to Rule 18.2, pointer subtraction between two pointers is allowed.

Example

```
void fn1 ( void )
{
    uint8_t  a[ 10 ];
    uint8_t  *ptr;
    uint8_t  index = 0U;

    index = index + 1U;      /* Compliant - rule only applies to pointers */

    a[ index ] = 0U;        /* Compliant */
    ptr = &a[ 5 ];          /* Compliant */

    ptr = a;
    ptr++;                  /* Compliant - increment operator not + */
    *( ptr + 5 ) = 0U;      /* Non-compliant */
    ptr[ 5 ] = 0U;          /* Compliant */
}
```

```

void fn2 ( void )
{
  uint8_t array_2_2[ 2 ][ 2 ] = { { 1U, 2U }, { 4U, 5U } };
  uint8_t i = 0U;
  uint8_t j = 0U;
  uint8_t sum = 0U;

  for ( i = 0U; i < 2U; i++ )
  {
    uint8_t *row = array_2_2[ i ];

    for ( j = 0U; j < 2U; j++ )
    {
      sum += row[ j ]; /* Compliant */
    }
  }
}

```

In the following example, Rule 18.1 may also be violated if `p1` does not point to an array with at least six elements and `p2` does not point to an array with at least 4 elements.

```

void fn3 ( uint8_t *p1, uint8_t p2[ ] )
{
  p1++; /* Compliant */
  p1 = p1 + 5; /* Non-compliant */
  p1[ 5 ] = 0U; /* Compliant */

  p2++; /* Compliant */
  p2 = p2 + 3; /* Non-compliant */
  p2[ 3 ] = 0U; /* Compliant */
}

uint8_t a1[ 16 ];
uint8_t a2[ 16 ];
uint8_t data = 0U;

void fn4 ( void )
{
  fn3 ( a1, a2 );
  fn3 ( &data, &a2[ 4 ] );
}

```

See also

Rule 18.1, Rule 18.2

Rule 18.5 Declarations should contain no more than two levels of pointer nesting

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

No more than two *pointer declarators* should be applied consecutively to a type. Any *typedef-name* appearing in a declaration is treated as if it were replaced by the type that it denotes.

Rationale

The use of more than two levels of pointer nesting can impair the ability to understand the behaviour of the code, and should therefore be avoided.

Example

```
typedef int8_t * INTPTR;

void function ( int8_t ** arrPar[ ] ) /* Non-compliant */
{
    int8_t ** obj2; /* Compliant */
    int8_t *** obj3; /* Non-compliant */
    INTPTR * obj4; /* Compliant */
    INTPTR * const * const obj5; /* Non-compliant */
    int8_t ** arr[ 10 ]; /* Compliant */
    int8_t ** ( *parr )[ 10 ]; /* Compliant */
    int8_t * ( **pparr )[ 10 ]; /* Compliant */
}

struct s
{
    int8_t * s1; /* Compliant */
    int8_t ** s2; /* Compliant */
    int8_t *** s3; /* Non-compliant */
};

struct s * ps1; /* Compliant */
struct s ** ps2; /* Compliant */
struct s *** ps3; /* Non-compliant */

int8_t ** ( *pfunc1 )( void ); /* Compliant */
int8_t ** ( **pfunc2 )( void ); /* Compliant */
int8_t ** ( ***pfunc3 )( void ); /* Non-compliant */
int8_t *** ( **pfunc4 )( void ); /* Non-compliant */
```

Note:

- `arrPar` is of type pointer to pointer to pointer to `int8_t` because parameters declared with array type are converted to a pointer to the initial element of the array — this is three levels and is non-compliant;
- `arr` is of type array of pointer to pointer to `int8_t` — this is compliant;
- `parr` is of type pointer to array of pointer to pointer to `int8_t` — this is compliant;
- `pparr` is of type pointer to pointer to array of pointer to `int8_t` — this is compliant.

Rule 18.6 The address of an object with automatic or thread-local storage shall not be copied to another object that persists after the first object has ceased to exist

C90 [Undefined 9, 26], C99 [Undefined 8, 9, 40], C11 [Undefined 9, 10, 43]

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

The address of an object might be copied by means of:

- *Assignment*;
- Memory move or copying functions.

Rationale

The address of an object becomes indeterminate when the lifetime of that object expires. Any use of an indeterminate address results in undefined behaviour.

Example

```
int8_t *func ( void )
{
    int8_t local_auto;

    return &local_auto;    /* Non-compliant - &local_auto is indeterminate
                          *                          when func returns          */
}
```

In the following example, the function `g` stores a copy of its pointer parameter `p`. If `p` **always** points to an object with static storage duration then the code is compliant with this rule. However, in the example given, `p` does point to an object with automatic storage duration. In such a case, copying the parameter `p` is non-compliant.

```
uint16_t *sp;

void g ( uint16_t *p )
{
    sp = p;                /* Non-compliant - address of f's parameter u
                          *                          copied to static sp          */
}

void f ( uint16_t u )
{
    g ( &u );
}

void h ( void )
{
    static uint16_t *q;

    uint16_t x = 0u;

    q = &x;                /* Non-compliant - &x stored in object with
                          *                          greater lifetime          */
}
```

Rule 18.7 Flexible array members shall not be declared

C99 [Undefined 59], C11 [Undefined 62]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C99, C11

Rationale

Flexible array members are most likely to be used in conjunction with dynamic memory allocation which is banned by Dir 4.12 and Rule 21.3.

The presence of flexible array members modifies the behaviour of the `sizeof` operator in ways that might not be expected by a programmer. The assignment of a structure that contains a flexible array member to another structure of the same type may not behave in the expected manner as it copies only those elements up to but not including the start of the flexible array member.

Example

```
#include <stdlib.h>

struct s
{
    uint16_t len;
    uint32_t data[ ]; /* Non-compliant - flexible array member */
} str;

struct s *copy ( struct s *s1 )
{
    struct s *s2;

    /* Omit malloc ( ) return check for brevity */
    s2 = malloc ( sizeof ( struct s ) + ( s1->len * sizeof ( uint32_t ) ) );

    *s2 = *s1;          /* Only copies s1->len */

    return s2;
}
```

See also

Dir 4.12, Rule 21.3

Rule 18.8 Variable-length arrays shall not be used

C99 [Unspecified 21; Undefined 69, 70], C11 [Unspecified 22; Undefined 16, 75, 76]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C99, C11

Rationale

Variable-length arrays are specified when the size of an array declared in a block or a function prototype is not an *integer constant expression*. They are typically implemented as a variable size object stored on the stack. Their use can therefore make it impossible to determine statically the amount of memory that must be reserved for a stack.

If the size of a variable-length array is negative or zero, the behaviour is undefined.

If a variable-length array is used in a context in which its type is required to be compatible with the type of another array, then the size of the array types shall be identical. Further, all sizes shall evaluate to positive integers. If these requirements are not met, the behaviour is undefined.

If a variable-length array is used in the operand of a *sizeof* operator, under some circumstances it is unspecified whether the array size expression is evaluated or not.

Each instance of a variable-length array type has its size fixed at the start of its lifetime. This gives rise to behaviour that might be confusing, for example:

```
void f ( void )
{
    uint16_t n = 5;

    typedef uint16_t Vector[ n ]; /* An array type with 5 elements */

    n = 7;

    Vector    a1;                /* An array type with 5 elements */

    uint16_t a2[ n ];           /* An array type with 7 elements */
}
```

Example

There is no use of variable-length arrays that is compliant with this rule. The examples show some of the undefined behaviour that can arise from their use.

```
void f ( int16_t n )
{
    uint16_t vla[ n ];          /* Non-compliant - Undefined if n <= 0 */
}

void g ( void )
{
    f ( 0 );                   /* Undefined */
    f ( -1 );                  /* Undefined */
    f ( 10 );                  /* Defined */
}
```

See also

Rule 13.6, Rule 18.10

Rule 18.9 An object with *temporary lifetime* shall not undergo array-to-pointer conversion

C99 [Undefined 8, 35], C11 [Undefined 9]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

Temporary lifetime is a storage duration which describes the lifetime of the elements of non-*lvalue* arrays.

An array which is a member of a non-*lvalue* expression with structure or union type shall not be used as a value, other than as the immediate *postfix-expression* operand to a subscript operator.

The subscript operator shall not be used to produce a modifiable *lvalue*.

Rationale

An array object can be a member of a structure or union, and therefore form part of the result value of any value expression. Because arrays used in an expression are always value-converted to a pointer to their elements, it is possible to form a pointer to such array sub-objects even when they are not part of a declared object with normal lifetime.

Modifying elements of such an array implicitly results in undefined behaviour in C90, and explicitly results in undefined behaviour in C99 and later. Accessing the elements of such an array after the end of its lifetime results in undefined behaviour, which is implicitly limited to the next sequence point in C90 and C99 (meaning the pointer cannot be stored or passed to any function), and to the duration of the complete containing expression in C11 and later.

The pointer will not be to `const`-qualified elements unless the array element type is `const`-qualified; therefore the array's type may appear to allow modification, and the generation of apparently modifiable element *lvalues*, even though the array itself has *temporary lifetime*.

Since the containing object can always be assigned to a named intermediate object by value, there is little reason to ever attempt to take the address of the value with temporary lifetime.

Example

```

/* Value object containing an array as an element */

struct S1 {
    int32_t array[10];
};

struct S1 s1;
struct S1 getS1 (void);

void foo(int32_t const * p);

/* Compliant - not temporary storage duration */

int32_t * p = s1.array;
s1.array[0] = 1;
foo( s1.array );

/* Non-compliant - temporary storage duration */

p = getS1().array;           /* also creates dangling pointer */
foo( getS1().array );
foo( (s1 = s1).array );     /* other forms of non-lvalue expression */

/* Compliant - immediate element access is always safe */

int32_t j = getS1().array[3]; /* element copied: const access */
j = (s1 = s1).array[3];

/* Non-compliant - element used as a modifiable lvalue */

getS1().array[3] = j;
(1 ? s1 : s1).array[3] = j;

```

Rule 18.10 Pointers to variably-modified array types shall not be used

C99 [Undefined 69, 70], C11 [Undefined 75, 76]

Category	Mandatory
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Amplification

A pointer to a variably-modified array type shall not be used in the declaration of any object or parameter.

A parameter declared to have an array type is not a pointer-to-array type (unless it is an array of arrays), because it is rewritten to a pointer to the element type.

Rationale

Compatibility between array types requires the size specifiers for the pointed-to arrays to have equal values. However, for variably-modified array types this cannot be determined at compile-time.

If two pointers to array types are used in any way that requires them to be compatible (such as assignment), and the size specifiers for the pointed-to array are not the same, the behaviour is undefined. This is undecidable in general, effectively leaving *all* such operations untyped.

Example

```
/* Non-compliant */
void f1 (uint16_t n, uint16_t (* a) [n])
{
    uint16_t ( *p ) [ 20 ];
    p = a; /* undefined unless n == 20, but types always assumed compatible */
}

/* Compliant */
void f2 (uint16_t n, uint16_t a[n])
{
    uint16_t * p;
    p = a; /* pointed-to type is not variably-modified, always well-defined */
}
```

See also

Rule 18.8

8.19 Overlapping storage

Rule 19.1 An object shall not be assigned or copied to an overlapping object

C90 [Undefined 34, 55], C99 [Undefined 51, 94], C11 [Undefined 54, 100]

Category Mandatory

Analysis Undecidable, System

Applies to C90, C99, C11

Rationale

The behaviour is undefined when two objects are created which have some overlap in memory and one is assigned or copied to the other.

Exception

The following are permitted because the behaviour is well-defined:

1. Assignment between two objects that overlap exactly and have compatible types (ignoring their type qualifiers)
2. Copying between objects that overlap partially or completely using the Standard Library function `memcpy`

Example

This example also violates Rule 19.2 because it uses unions.

```
void fn ( void )
{
    union
    {
        int16_t i;
        int32_t j;
    } a = { 0 };

    a.j = a.i;    /* Non-compliant */
}

#include <string.h>

int16_t a[ 20 ];

void f ( void )
{
    memcpy ( &a[ 5 ], &a[ 4 ], 2u * sizeof ( a[ 0 ] ) ); /* Non-compliant */
}

void g ( void )
{
    int16_t *p = &a[ 0 ];
    int16_t *q = &a[ 0 ];

    *p = *q;    /* Compliant - exception 1 */
}
```

See also

Rule 19.2

Rule 19.2 The *union* keyword should not be used

C90 [Undefined 39, 40; Implementation 27]

C99 [Unspecified 10; Undefined 61, 62]

C11 [Unspecified 11; Undefined 64, 65]

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

A union member can be written and the same member can then be read back in a well-defined manner.

However, if a union member is written and then a different union member is read back, the behaviour depends on the relative sizes of the members:

- If the member read is wider than the member written then the value is unspecified;
- Otherwise, the value is implementation-defined.

The C Standard permits the bytes of a union member to be accessed by means of another member whose type is array of *unsigned char*. However, since it is possible to access bytes with unspecified values, unions should not be used.

If this rule is not followed, the kinds of behaviour that need to be determined are:

- Padding — how much padding is inserted at the end of the union;
- Alignment — how are members of any structures within the union aligned;
- Endianness — is the most significant byte of a word stored at the lowest or highest memory address;
- Bit-order — how are bits numbered within bytes and how are bits allocated to bit fields.

Example

In this non-compliant example, a 16-bit value is stored into a union but a 32-bit value is read back resulting in an unspecified value being returned.

```
uint32_t zext ( uint16_t s )
{
    union
    {
        uint32_t ul;
        uint16_t us;
    } tmp;

    tmp.us = s;
    return tmp.ul; /* unspecified value */
}
```

See also

Rule 19.1

8.20 Preprocessing directives

Rule 20.1 *#include* directives should only be preceded by preprocessor directives or comments

C90 [Undefined 56], C99 [Undefined 96, 97], C11 [Undefined 102, 103]

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

The rule shall be applied to the contents of a file before preprocessing occurs.

Rationale

To aid code readability, all the *#include* directives in a particular code file should be grouped together near the top of the file.

Additionally, using *#include* to include a standard *header file* within a declaration or definition, or using part of the Standard Library before the inclusion of the related standard *header file* leads to undefined behaviour.

Example

```
/* f.h */
xyz = 0;

/* f.c */

int16_t
#include "f.h"    /* Non-compliant */

/* f1.c */
#define F1_MACRO

#include "f1.h"   /* Compliant */
#include "f2.h"   /* Compliant */

int32_t i = 0;

#include "f3.h"   /* Non-compliant */
```

Rule 20.2 The `'`, `"` or `\` characters and the `/*` or `//` character sequences shall not occur in a *header file* name

C90 [Undefined 14], C99 [Undefined 31], C11 [Undefined 34]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

The behaviour is undefined if:

- The `'`, `"` or `\` characters, or the `/*` or `//` character sequences are used between `<` and `>` delimiters in a header name preprocessing token;
- The `'` or `\` characters, or the `/*` or `//` character sequences are used between the `"` delimiters in a header name preprocessing token.

Note: although use of the `\` character results in undefined behaviour, many implementations will accept the `/` character in its place.

Example

```
#include "fi'le.h" /* Non-compliant */
```

Rule 20.3 The `#include` directive shall be followed by either a `<filename>` or `"filename"` sequence

C90 [Undefined 48], C99 [Undefined 85], C11 [Undefined 91]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule applies after macro replacement has been performed.

Rationale

The behaviour is undefined if a `#include` directive does not use one of the following forms:

- `#include <filename>`
- `#include "filename"`

Example

```
#include "filename.h" /* Compliant */
#include <filename.h> /* Compliant */
#include another.h /* Non-compliant */
```

```
#define HEADER "filename.h"
#include HEADER          /* Compliant */
#define FILENAME file2.h
#include FILENAME        /* Non-compliant */

#define BASE "base"
#define EXT ".ext"
#include BASE EXT        /* Non-compliant - strings are concatenated
                        * after preprocessing */

#include "../include/cpu.h" /* Compliant - filename may include a path */
```

Rule 20.4 A macro shall not be defined with the same name as a keyword

C90 [Undefined 56], C99 [Undefined 98], C11 [Undefined 104]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

This rule applies to all keywords, including those that implement language extensions.

Rationale

Using macros to change the meaning of keywords can be confusing. The behaviour is undefined if a standard header is included while a macro is defined with the same name as a keyword.

Example

The following example is non-compliant because it alters the behaviour of the *int* keyword. Including a standard header in the presence of this macro results in undefined behaviour.

```
#define int some_other_type
#include <stdlib.h>
```

The following example shows that it is non-compliant to redefine the keyword *while* but it is compliant to define a macro that expands to statements.

```
#define while( E ) for ( ; ( E ) ; ) /* Non-compliant - redefined while */
#define unless( E ) if ( ! ( E ) ) /* Compliant */

#define seq( S1, S2 ) do { \
    S1; S2; } while ( false ) /* Compliant */
#define compound( S ) { S; } /* Compliant */
```

The following example is only compliant with C90.

```
/* Remove inline if compiling for C90 */
#define inline
```

See also

Rule 21.1

Rule 20.5 `#undef` should not be used

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

The use of `#undef` can make it unclear which macros exist at a particular point within a translation unit.

Example

```
#define QUALIFIER volatile

#undef QUALIFIER          /* Non-compliant */

void f ( QUALIFIER int32_t p )
{
    while ( p != 0 )
    {
        ;                /* Wait... */
    }
}
```

Rule 20.6 Tokens that look like a preprocessing directive shall not occur within a macro argument

C90 [Undefined 50], C99 [Undefined 87], C11 [Undefined 93]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

An argument containing sequences of tokens that would otherwise act as preprocessing directives leads to undefined behaviour.

Example

```
#define M( A ) printf ( #A )

#include <stdio.h>

void main ( void )
{
    M (
#ifdef SW          /* Non-compliant */
    "Message 1"
#else             /* Non-compliant */
    "Message 2"
#endif
    );
}
```

The above may print

```
#ifdef SW "Message 1" #else "Message 2" #endif
```

or

```
"Message 2"
```

or exhibit some other behaviour.

Rule 20.7 Expressions resulting from the expansion of macro parameters shall be enclosed in parentheses

[Koenig 78-81]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

If the expansion of any macro parameter produces a token, or sequence of tokens, that form an expression then that expression, in the fully-expanded macro, shall either:

- Be a parenthesized expression itself; or
- Be enclosed in parentheses.

Note: this does not necessarily require that all macro parameters are parenthesized; it is acceptable for parentheses to be provided in macro arguments.

Rationale

If parentheses are not used, then operator precedence may not give the desired results when macro substitution occurs.

If a macro parameter is not being used as an expression then the parentheses are not necessary because no operators are involved.

Example

In the following non-compliant example,

```
#define M1( x, y ) ( x * y )
```

```
r = M1 ( 1 + 2, 3 + 4 );
```

the macro expands to give:

```
r = ( 1 + 2 * 3 + 4 );
```

The expressions $1 + 2$ and $3 + 4$ are derived from expansion of parameters x and y respectively, but neither is enclosed in parentheses. The value of the resulting expression is 11, whereas the result 21 might have been expected.

The code could be written in a compliant manner either by parenthesizing the macro arguments, or by writing an alternative version of the macro that inserts parentheses during expansion, for example:

```
r = M1 ( ( 1 + 2 ), ( 3 + 4 ) ); /* Compliant */
#define M2( x, y ) ( ( x ) * ( y ) )
r = M2 ( 1 + 2, 3 + 4 ); /* Compliant */
```

The following example is compliant because the first expansion of `x` is as the operand of the `##` operator, which does not produce an expression. The second expansion of `x` is as an expression which is parenthesized as required.

```
#define M3( x ) a ## x = ( x )
int16_t M3 ( 0 );
```

The following example is compliant because expansion of the parameter `M` as a member name does not produce an expression. Expansion of the parameter `S` produces an expression, with structure or union type, which does require parentheses.

```
#define GET_MEMBER( S, M ) ( S ).M
v = GET_MEMBER ( s1, minval );
```

The following compliant example shows that it is not always necessary to parenthesize every instance of a parameter, although this is often the easiest method of complying with this rule.

```
#define F( X ) G( X )
#define G( Y ) ( ( Y ) + 1 )
int16_t x = F ( 2 );
```

The fully-expanded macro is `((2) + 1)`. Tracing back through macro expansion, the value 2 arises from expansion of parameter `Y` in macro `G` which in turn arises from parameter `x` in macro `F`. Since 2 is parenthesized in the fully-expanded macro, the code is compliant.

See also

Dir 4.9

Rule 20.8 The controlling expression of a `#if` or `#elif` preprocessing directive shall evaluate to 0 or 1

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule does not apply to controlling expressions in preprocessing directives which are not evaluated. Controlling expressions are not evaluated if they are within code that is being excluded and cannot have an effect on whether code is excluded or not.

Rationale

Strong typing requires the controlling expression of conditional inclusion preprocessing directives to have a Boolean value.

Example

```
#define FALSE 0
#define TRUE 1

#if FALSE          /* Compliant          */
#endif

#if 10             /* Non-compliant        */
#endif

#if ! defined ( X ) /* Compliant          */
#endif

#if A > B          /* Compliant assuming A and B are numeric */
#endif
```

See also

Rule 14.4

Rule 20.9 All identifiers used in the controlling expression of *#if* or *#elif* preprocessing directives shall be *#define*'d before evaluation

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

As well as using a *#define* preprocessor directive, identifiers may effectively be *#define*'d in other, implementation-defined, ways. For example some implementations support:

- Using a compiler command-line option, such as `-D` to allow identifiers to be defined prior to translation;
- Using environment variables to achieve the same effect;
- Pre-defined identifiers provided by the compiler.

Rationale

If an attempt is made to use a macro identifier in a preprocessor directive, and that identifier has not been defined, then the preprocessor will assume that it has a value of zero. This may not meet developer expectations.

Example

The following examples assume that the macro `M` is undefined.

```
#if M == 0          /* Non-compliant          */
                    /* Does 'M' expand to zero or is it undefined? */
#endif

#if defined ( M )  /* Compliant - M is not evaluated          */
#if M == 0         /* Compliant - M is known to be defined    */
                    /* 'M' must expand to zero.                */
#endif
#endif
```

```
/* Compliant - B is only evaluated in ( B == 0 ) if it is defined */
#if defined ( B ) && ( B == 0 )
#endif
```

Rule 20.10 The # and ## preprocessor operators should not be used

C90 [Unspecified 12; Undefined 51, 52]
C99 [Unspecified 25; Undefined 3, 88, 89]
C11 [Unspecified 26; Undefined 3, 94, 95]

Category Advisory
Analysis Decidable, Single Translation Unit
Applies to C90, C99, C11

Rationale

The order of evaluation associated with multiple #, multiple ## or a mix of # and ## preprocessor operators is unspecified. In some cases it is therefore not possible to predict the result of macro expansion.

The use of the ## operator can result in code that is obscure.

Note: Rule 1.3 covers the undefined behaviour that arises if either:

- The result of a # operator is not a valid string literal; or
- The result of a ## operator is not a valid preprocessing token.

See also

Rule 20.11

Rule 20.11 A macro parameter immediately following a # operator shall not immediately be followed by a ## operator

C90 [Unspecified 12], C99 [Unspecified 25], C11 [Unspecified 26]

Category Required
Analysis Decidable, Single Translation Unit
Applies to C90, C99, C11

Rationale

The order of evaluation associated with multiple #, multiple ## or a mix of # and ## preprocessor operators is unspecified. The use of # and ## is discouraged by Rule 20.10. In particular, the result of a # operator is a string literal and it is extremely unlikely that pasting this to any other preprocessing token will result in a valid token.

Example

```
#define A( x )      #x          /* Compliant */
#define B( x, y )  x ## y      /* Compliant */
#define C( x, y )  #x ## y     /* Non-compliant */
```

See also

Rule 20.10

Rule 20.12 A macro parameter used as an operand to the # or ## operators, which is itself subject to further macro replacement, shall only be used as an operand to these operators

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Rationale

A macro parameter that is used as an operand of a # or ## operator is **not** expanded prior to being used. The same parameter appearing elsewhere in the replacement text **is** expanded. If the macro parameter is itself subject to macro replacement, its use in mixed contexts within a macro replacement may not meet developer expectations.

Example

In the following non-compliant example, the macro parameter `x` is replaced with `AA` which is subject to further macro replacement when not used as the operand of `##`.

```
#define AA          0xffff
#define BB( x )    ( x ) + wow ## x /* Non-compliant */

void f ( void )
{
    int32_t wowAA = 0;

    /* Expands as wowAA = ( 0xffff ) + wowAA; */
    wowAA = BB ( AA );
}
```

In the following compliant example, the macro parameter `x` is not subject to further macro replacement.

```
int32_t speed;
int32_t speed_scale;
int32_t scaled_speed;

#define SCALE( X ) ( ( X ) * X ## _scale )

/* expands to scaled_speed = ( ( speed ) * speed_scale ); */
scaled_speed = SCALE ( speed );
```

Rule 20.13 A line whose first token is # shall be a valid preprocessing directive

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

White-space is permitted between the # and preprocessing tokens.

Rationale

A preprocessor directive may be used to conditionally exclude source code until a corresponding *#else*, *#elif* or *#endif* directive is encountered. A malformed or invalid preprocessing directive contained within the excluded source code may not be detected by the compiler, possibly leading to the exclusion of more code than was intended.

Requiring all preprocessor directives to be syntactically valid, even when they occur within an excluded block of code, ensures that this cannot happen.

Example

In the following example all the code between the *#ifndef* and *#endif* directives may be excluded if *AAA* is defined. The developer intended that *AAA* be assigned to *x*, but the *#else* directive was entered incorrectly and not diagnosed by the compiler.

```
#define AAA 2

int32_t foo ( void )
{
    int32_t x = 0;

#ifndef AAA
    x = 1;
#else1          /* Non-compliant */
    x = AAA;
#endif

    return x;
}
```

The following example is compliant because the text *#start* appearing in a comment is not a token.

```
/*
#start is not a token in a comment
*/
```

Rule 20.14 All `#else`, `#elif` and `#endif` preprocessor directives shall reside in the same file as the `#if`, `#ifdef` or `#ifndef` directive to which they are related

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

Confusion can arise when blocks of code are included or excluded by the use of conditional compilation directives which are spread over multiple files. Requiring that a `#if` directive be terminated within the same file reduces the visual complexity of the code and the chance that errors will be made during maintenance.

Note: `#if` directives may be used within included files provided they are terminated within the same file.

Example

```
/* file1.c */
#ifdef A          /* Compliant */
#include "file1.h"
#endif
/* End of file1.c */

/* file2.c */
#if 1            /* Non-compliant */
#include "file2.h"
/* End of file2.c*/

/* file1.h */
#if 1           /* Compliant */
#endif
/* End of file1.h */

/* file2.h */
#endif
/* End of file2.h */
```

8.21 Standard libraries

Rule 21.1 *#define* and *#undef* shall not be used on a reserved identifier or reserved macro name

C90 [Undefined 54, 57, 58, 62, 71]

C99 [Undefined 93, 100, 101, 104, 108, 116, 118, 130]

C11 [Undefined 99, 106, 107, 110, 114, 122, 126, 138]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

This rule applies to the following:

- Identifiers or macro names beginning with an underscore;
- Identifiers in file scope described in Section 7, “Library”, of the C Standard;
- Macro names described in Section 7, “Library”, of the C Standard as being defined in a standard header.

This rule also prohibits the use of *#define* or *#undef* on the identifier *defined* as this results in explicitly undefined behaviour.

This rule does **not** include those identifiers or macro names that are described in the section of the applicable C Standard entitled “Future Library Directions”.

The C Standard states that defining a macro with the same name as:

- A macro defined in a standard header, or
- An identifier with file scope declared in a standard header

is well-defined provided that the header is not included. This rule does not permit such definitions on the grounds that they are likely to cause confusion.

Note: the macro `NDEBUG` is not defined in a standard header and may therefore be *#define*'d.

Rationale

Reserved identifiers and reserved macro names are intended for use by the implementation. Removing or changing the meaning of a reserved macro may result in undefined behaviour.

Example

```
#undef  __LINE__          /* Non-compliant - begins with _ */
#define  __GUARD_H 1      /* Non-compliant - begins with _ */
#undef   __BUILTIN_sqrt   /* Non-compliant - the implementation
                          * may use __BUILTIN_sqrt for other
                          * purposes, e.g. generating a sqrt
                          * instruction */
```

```
#define defined                /* Non-compliant - reserved identifier */
#define errno my_errno        /* Non-compliant - library identifier */
#define isneg( x ) ( ( x ) < 0 ) /* Compliant - rule doesn't include
*                               future library
*                               directions */
```

See also

Rule 20.4, Rule 20.5

Rule 21.2 A reserved identifier or reserved macro name shall not be declared

C90 [Undefined 57, 58, 64, 71]

C99 [Undefined 93, 100, 101, 104, 108, 116, 118, 130]

C11 [Undefined 99, 106, 107, 110, 114, 122, 126, 138]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

See the Amplification for Rule 21.1 for a description of the relevant identifiers and macro names.

Rationale

The implementation is permitted to rely on reserved identifiers behaving as described in the C Standard and may treat them specially. If reserved identifiers are reused, the program may exhibit undefined behaviour.

Example

In the following non-compliant example, the function *memcpy* is declared explicitly. The compliant method of declaring this function is to include `<string.h>`.

```
/*
 * Include <stddef.h> to define size_t
 */
#include <stddef.h>
extern void *memcpy ( void *restrict s1, const void *restrict s2, size_t n );
```

An implementation is permitted to provide a *function-like macro* definition for each Standard Library function **in addition to** the library function itself. This feature is often used by compiler writers to generate efficient inline operations in place of the call to a library function. Using a *function-like macro*, the call to a library function can be replaced with a call to a reserved function that is detected by the compiler's code generation phase and replaced with the inline operation. For example, the fragment of `<math.h>` that declares *sqrt* might be written using a *function-like macro* that generates a call to `_BUILTIN_sqrt` which is replaced with an inline `SQRT` instruction on processors that support it:

```
extern double sqrt ( double x );

#define sqrt( x ) ( _BUILTIN_sqrt ( x ) )
```

The following non-compliant code might interfere with the compiler's built-in mechanism for handling `sqrt` and therefore produce undefined behaviour:

```
static double _BUILTIN_sqrt ( double x )      /* Non-compliant          */
{
    return x * x;
}

#include <math.h>

float64_t x = sqrt ( ( float64_t ) 2.0 );    /* sqrt may not behave as
                                             * defined in the C Standard */
```

Rule 21.3 The memory allocation and deallocation functions of `<stdlib.h>` shall not be used

C90 [Unspecified 19; Undefined 9, 91, 92; Implementation 69]
 C99 [Unspecified 39, 40; Undefined 8, 9, 168–171; Implementation J.3.12(35)]
 C11 [Unspecified 42, 43; Undefined 9, 10, 177–181; Implementation J.3.12(37)]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The identifiers *calloc*, *malloc*, *realloc*, *aligned_alloc* and *free* shall not be used and no macro with one of these names shall be expanded.

Rationale

Use of dynamic memory allocation and deallocation routines provided by the Standard Library can lead to undefined behaviour, for example:

- Memory that was not dynamically allocated is subsequently freed;
- A pointer to freed memory is used in any way;
- Accessing allocated memory before storing a value into it.

Note: this rule is a specific instance of Dir 4.12.

See also

Dir 4.12, Rule 18.7, Rule 22.1, Rule 22.2

Rule 21.4 The standard *header file* `<setjmp.h>` shall not be used

C90 [Unspecified 14; Undefined 64–67]
 C99 [Unspecified 32; Undefined 118–121, 173]
 C11 [Unspecified 35; Undefined 124–127, 183]
 [Koenig 74]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The standard *header file* `<setjmp.h>` shall not be *#include'd*, and none of the features that are specified as being provided by `<setjmp.h>` shall be used.

Rationale

setjmp and *longjmp* allow the normal function call mechanisms to be bypassed. Their use may lead to undefined and unspecified behaviour.

Rule 21.5 The standard *header file* `<signal.h>` shall not be used

C90 [Undefined 67–69; Implementation 48–52]
 C99 [Undefined 122–127; Implementation J.3.12(12)]
 C11 [Undefined 128–135; Implementation J.3.12(14)]
 [Koenig 74]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The standard *header file* `<signal.h>` shall not be *#include'd*, and none of the features that are specified as being provided by `<signal.h>` shall be used.

Rationale

Signal handling contains implementation-defined and undefined behaviour.

Rule 21.6 The Standard Library input/output functions shall not be used

C90 [Unspecified 2–5, 16–18; Undefined 77–89; Implementation 53–68]
 C99 [Unspecified 3–6, 34–37; Undefined 138–166, 186; Implementation J.3.12(14–32)]
 C11 [Unspecified 4–7, 37–40; Undefined 146–175, 198; Implementation J.3.12(16–34)]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule applies to the functions that are specified as being provided by `<stdio.h>` and the wide-character equivalents specified as being provided by `<wchar.h>`.

None of these identifiers shall be used and no macro with one of these names shall be expanded.

Rationale

Streams and file I/O have unspecified, undefined and implementation-defined behaviours associated with them.

See also

Rule 22.1, Rule 22.3, Rule 22.4, Rule 22.5, Rule 22.6

Rule 21.7 The Standard Library functions *atof*, *atoi*, *atol* and *atoll* of `<stdlib.h>` shall not be used

C90 [Undefined 90], C99 [Undefined 113], C11 [Undefined 119]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The identifiers *atof*, *atoi*, *atol* and *atoll* shall not be used and no macro with one of these names shall be expanded.

Rationale

These functions have undefined behaviour associated with them when the string cannot be converted.

Rule 21.8 The Standard Library termination functions of `<stdlib.h>` shall not be used

C90 [Undefined 93; Implementation 70–71]
 C99 [Undefined 172; Implementation J.3.12(36–37)]
 C11 [Undefined 182, 185; Implementation J.3.12(38–39)]

Category Required
Analysis Decidable, Single Translation Unit
Applies to C90, C99, C11

Amplification

The termination functions are *abort*, *exit*, *_Exit* and *quick_exit*.

Identifiers with these names shall not be used and no macro with one of these names shall be expanded.

Rationale

These functions have undefined and implementation-defined behaviours associated with them.

Rule 21.9 The Standard Library functions *bsearch* and *qsort* of `<stdlib.h>` shall not be used

C90 [Unspecified 20, 21]
 C99 [Unspecified 41, 42; Undefined 176–178]
 C11 [Unspecified 46, 47; Undefined 187–189]

Category Required
Analysis Decidable, Single Translation Unit
Applies to C90, C99, C11

Amplification

The identifiers *bsearch* and *qsort* shall not be used and no macro with one of these names shall be expanded.

Rationale

If the comparison function does not behave consistently when comparing elements, or it modifies any of the elements, the behaviour is undefined.

Note: the unspecified behaviour, which relates to the treatment of elements that compare as equal, can be avoided by ensuring that the comparison function never returns 0. When two elements are otherwise equal, the comparison function could return a value that indicates their relative order in the initial array.

The implementation of *qsort* is likely to be recursive and will therefore place unknown demands on stack resource. This is of concern in embedded systems as the stack is likely to be a fixed, often small, size.

Rule 21.10 The Standard Library time and date functions shall not be used

C90 [Unspecified 22; Undefined 80, 97; Implementation 75, 76]
 C99 [Unspecified 43, 44; Undefined 146, 154, 182; Implementation J.3.12(39-42)]
 C11 [Unspecified 48, 49; Undefined 154, 162, 193, 197; Implementation J.3.12(41-45)]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

This rule applies to the functions that are specified as being provided by `<time.h>` and the function `wcsftime` provided by `<wchar.h>`.

The standard *header file* `<time.h>` shall not be `#include'd`, and none of the features that are specified as being provided by `<time.h>` shall be used.

For C99 or later, the function `wcsftime` shall not be used and no macro with this name shall be expanded.

Rationale

The time and date functions have unspecified, undefined and implementation-defined behaviours associated with them.

Rule 21.11 The standard *header file* `<tgmath.h>` should not be used

C99 [Undefined 184, 185], C11 [Undefined 195, 196]

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Amplification

The standard *header file* `<tgmath.h>` should not be `#include'd`.

Note: Due to the duplication of macro names between `<tgmath.h>`, `<math.h>` and `<complex.h>` this rule does not have the additional requirement that *none of the features that are specified as being provided by `<tgmath.h>` should be used*, as use by means of `#including` either of these other standard *header files* is not constrained. Any other definition of a macro specified as being provided by `<tgmath.h>` will be a violation of Rule 21.1 and/or Rule 21.2.

Rationale

Using the facilities of `<tgmath.h>` may result in undefined behaviour.

Example

```
#include <tgmath.h>

float f1, f2;

void f ( void )
{
    f1 = sqrt ( f2 );    /* Non-compliant - generic sqrt used    */
}

#include <math.h>

float f1, f2;

void f ( void )
{
    f1 = sqrtf ( f2 );  /* Compliant - float version of sqrt used */
}
```

See also

Rule 21.1, Rule 21.2, Rule 21.22, Rule 21.23

Rule 21.12 The standard *header file* `<fenv.h>` shall not be used

C99 [Unspecified 27, 28; Undefined 109–112; Implementation J.3.6(8)]

C11 [Unspecified 27, 28; Undefined 115–118; Implementation J.3.6(9)]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C99, C11

Amplification

The standard *header file* `<fenv.h>` shall not be *#include'd*, and none of the features that are specified as being provided by `<fenv.h>` shall be used.

Rationale

In some circumstances, the values of the floating-point status flags are unspecified and attempts to access them may lead to undefined behaviour.

The order in which exceptions are raised by the *feraiseexcept* function is unspecified and could therefore result in a program that has been designed for a certain order not operating correctly.

Calling the *fesetenv* or *feupdateenv* functions with invalid arguments results in *undefined behaviour*.

Calling the `fesetround` function should be done with care because:

1. Setting the rounding mode may have unexpected consequences, e.g. setting the current rounding mode to upwards does not guarantee that the result of evaluating an expression is an upward approximation to the value of the expression over the reals.
2. Several implementations of functions declared in `<math.h>` have been designed to support round-to-nearest only: if such functions are called when a different rounding mode is set, the results can be unpredictable.

Note: In *conforming implementations*, the rounding direction mode is set to rounding to nearest at program start-up.

Example

```
#include <fenv.h>

void f ( float32_t x, float32_t y )
{
    float32_t z;

    feclearexcept ( FE_DIVBYZERO );          /* Non-compliant */

    z = x / y;

    if ( fetestexcept ( FE_DIVBYZERO ) )    /* Non-compliant */
    {
    }

    else
    {
#pragma STDC FENV_ACCESS ON

        z = x * y;
    }

    if ( z > x )
    {
#pragma STDC FENV_ACCESS OFF

        if ( fetestexcept ( FE_OVERFLOW ) ) /* Non-compliant */
        {
        }
    }
}
```

Rule 21.13 Any value passed to a function in `<ctype.h>` shall be representable as an *unsigned char* or be the value `EOF`

C90 [Undefined 63], C99 [Undefined 107], C11 [Undefined 113]

Category	Mandatory
Analysis	Undecidable, System
Applies to	C90, C99, C11

Rationale

The relevant functions from `<ctype.h>` are defined to take an *int* argument where the expected value is either in the range of an *unsigned char* or is a negative value equivalent to `EOF`. The use of any other values results in *undefined behaviour*.

Example

Note: The `int` casts in the following example are required to comply with Rule 10.3.

```
bool_t f ( uint8_t a )
{
    return (    isdigit ( ( int32_t ) a )           /* Compliant */
            && isalpha ( ( int32_t ) 'b' )        /* Compliant */
            && islower ( EOF )                   /* Compliant */
            && isalpha ( 256 ) );               /* Non-compliant */
}
```

See also

Rule 10.3

Rule 21.14 The Standard Library function *memcmp* shall not be used to compare null terminated strings

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

For the purposes of this rule, “null terminated strings” are:

- String literals;
- Arrays having *essentially character type* which contain a *null character*.

Rationale

The Standard Library function `int memcmp (const void *s1, const void *s2, size_t n);` performs a byte by byte comparison of the first `n` bytes of the two objects pointed at by `s1` and `s2`.

If *memcmp* is used to compare two strings and the length of either is less than `n`, then they may compare as different even when they are logically the same (i.e. each has the same sequence of characters before the null terminator) as the characters after a null terminator will be included in the comparison even though they do not form part of the string.

Example

```
extern char buffer1[ 12 ];
extern char buffer2[ 12 ];

void f1 ( void )
{
    ( void ) strcpy ( buffer1, "abc" );
    ( void ) strcpy ( buffer2, "abc" );

    /* The following use of memcmp is non-compliant */
    if ( memcmp ( buffer1, buffer2, sizeof ( buffer1 ) ) != 0 )
    {
        /*
         * The strings stored in buffer1 and buffer 2 are reported to be
         * different, but this may actually be due to differences in the
         * uninitialized characters stored after the null terminators.
         */
    }
}

/* The following definition violates other guidelines */
unsigned char headerStart[ 6 ] = { 'h', 'e', 'a', 'd', 0, 164 };

void f2 ( const uint8_t *packet )
{
    /* The following use of memcmp is compliant */
    if ( ( NULL != packet ) && ( memcmp( packet, headerStart, 6 ) == 0 ) )
    {
        /*
         * Comparison of values having essentially unsigned type reports that
         * contents are the same. Any null terminator is simply treated as a
         * zero value and any differences beyond it are significant.
         */
    }
}
```

See also

Rule 21.15, Rule 21.16

Rule 21.15 The pointer arguments to the Standard Library functions *memcpy*, *memmove* and *memcmp* shall be pointers to qualified or unqualified versions of compatible types

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

The Standard Library functions

```
void * memcpy ( void * restrict s1, const void * restrict s2, size_t n );
void * memmove ( void *s1, const void *s2, size_t n );
int memcmp ( const void *s1, const void *s2, size_t n );
```

perform a byte by byte copy, move or comparison of the first *n* bytes of the two objects pointed at by *s1* and *s2*.

An attempt to call one of these functions with arguments which are pointers to different types may indicate a mistake.

Example

```
/*
 * Is it intentional to only copy part of 's2'?
 */
void f ( uint8_t s1[ 8 ], uint16_t s2[ 8 ] )
{
    ( void ) memcpy ( s1, s2, 8 );    /* Non-compliant */
}
```

See also

Rule 21.14, Rule 21.16

Rule 21.16 The pointer arguments to the Standard Library function *memcpy* shall point to either a pointer type, an *essentially signed* type, an *essentially unsigned* type, an *essentially Boolean* type or an *essentially enum* type

C99 [Unspecified 9], C11 [Unspecified 10]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Rationale

The Standard Library function

```
int memcmp ( const void *s1, const void *s2, size_t n );
```

performs a byte by byte comparison of the first *n* bytes of the two objects pointed at by *s1* and *s2*.

Structures shall not be compared using *memcmp* as it may incorrectly indicate that two structures are not equal, even when their members hold the same values. Structures may contain padding with an indeterminate value between their members and *memcmp* will include this in its comparison. It cannot be assumed that the padding will be equal, even when the values of the structure members are the same. Unions have similar concerns along with the added complication that they may incorrectly be reported as having the same value when the representation of different, overlapping members are coincidentally the same.

Objects with *essentially floating* type shall not be compared with *memcmp* as the same value may be stored using different representations.

If an *essentially char* array contains a null character, it is possible to treat the data as a character string rather than simply an array of characters. However that distinction is a matter of interpretation rather than syntax. Since *essentially char* arrays are most frequently used to store character strings, an attempt to compare such arrays using *memcmp* (rather than *strcmp* or *strncmp*) may indicate an error as the number of characters to be compared will be determined by the value of the *size_t* argument rather than the location of the null characters used to terminate the strings. The result may therefore depend on the comparison of characters which are not part of the respective strings.

Example

```

struct S;

/*
 * Return value may indicate that 's1' and 's2' are different due to padding.
 */
bool_t f1 ( struct S *s1, struct S *s2 )
{
    return ( memcmp ( s1, s2, sizeof ( struct S ) ) != 0 );    /* Non-compliant */
}

union U
{
    uint32_t range;
    uint32_t height;
};

/*
 * Return value may indicate that 'u1' and 'u2' are the same
 * due to unintentional comparison of 'range' and 'height'.
 */
bool_t f2 ( union U *u1, union U *u2 )
{
    return ( memcmp ( u1, u2, sizeof ( union U ) ) != 0 );    /* Non-compliant */
}

const char a[ 6 ] = "task";

/*
 * Return value may incorrectly indicate strings are different as the
 * length of 'a' (4) is less than the number of bytes compared (6).
 */
bool_t f3 ( const char b[ 6 ] )
{
    return ( memcmp ( a, b, 6 ) != 0 );    /* Non-compliant */
}

```

See also

Rule 21.14, Rule 21.15

Rule 21.17 Use of the string handling functions from `<string.h>` shall not result in accesses beyond the bounds of the objects referenced by their pointer parameters

C90 [Undefined 96], C99 [Undefined 103, 180], C11 [Undefined 109, 191]

Category	Mandatory
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

The relevant string handling functions from `<string.h>` are:

`strcat`, `strchr`, `strcmp`, `strcoll`, `strcpy`, `strcspn`, `strlen`, `strpbrk`, `strrchr`, `strspn`, `strstr`, `strtok`

Rationale

Incorrect use of a function listed above may result in a read or write access beyond the bounds of an object passed as a parameter, resulting in *undefined behaviour*.

Example

```
char string[] = "Short";

void f1 ( const char *str )
{
    /*
     * Non-compliant use of strcpy as it results in writes beyond the end of 'string'
     */
    ( void ) strcpy ( string, "Too long to fit" );

    /*
     * Compliant use of strcpy as 'string' is only modified if 'str' will fit.
     */
    if ( strlen ( str ) < ( sizeof ( string ) - 1u ) )
    {
        ( void ) strcpy ( string, str );
    }
}

size_t f2 ( void )
{
    char text[ 5 ] = "Token";

    /*
     * The following is non-compliant as it results in reads beyond
     * the end of 'text' as there is no null terminator.
     */
    return strlen ( text );
}
```

See also

Rule 21.18

Rule 21.18 The `size_t` argument passed to any function in `<string.h>` shall have an appropriate value

C90 [Undefined 96], C99 [Undefined 103, 180, 181], C11 [Undefined 109, 191, 192]

Category Mandatory

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

The relevant functions in `<string.h>` are:

`memchr`, `memcmp`, `memcpy`, `memmove`, `memset`, `strncat`, `strncmp`, `strcpy`, `strxfrm`

An appropriate value is:

- Positive;
- No greater than the size of the smallest object passed to the function through a pointer parameter.

Rationale

Incorrect use of a function listed above may result in a read or write access beyond the bounds of an object passed as a parameter, resulting in *undefined behaviour*.

Example

```
char buf1[ 5 ] = "12345";
char buf2[ 10 ] = "1234567890";

void f ( void )
{
    if ( memcmp ( buf1, buf2, 5 ) == 0 )           /* Compliant */
    {
    }

    if ( memcmp ( buf1, buf2, 6 ) == 0 )           /* Non-compliant */
    {
    }
}
```

See also

Rule 21.17

Rule 21.19 The pointers returned by the Standard Library functions *localeconv*, *getenv*, *setlocale* or *strerror* shall only be used as if they have pointer to const-qualified type

C90 [Undefined], C99 [Undefined 114, 115, 174], C11 [Undefined 120, 121, 184]

Category Mandatory

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

The *localeconv* function returns a pointer of type `struct lconv *`. This pointer shall be regarded as if it had type `const struct lconv *`.

A `struct lconv` object includes pointers of type `char *` and the *getenv*, *setlocale*, and *strerror* functions each return a pointer of type `char *`. These pointers are used to access strings (null terminated arrays of type *char*). For the purpose of this rule, these pointers shall be regarded as if they had type `const char *`.

Rationale

The C Standard states that *undefined behaviour* occurs if a program modifies:

- The structure pointed to by the value returned by *localeconv*;
- The strings returned by *getenv*, *setlocale* or *strerror*.

Note: The C Standard does not specify the behaviour that results if the strings referenced by the structure pointed to by the value returned by *localeconv* are modified. This rule prohibits any changes to these strings as they are considered to be undesirable.

Treating the pointers returned by the various functions as if they were const-qualified allows an analysis tool to detect any attempt to modify an object through one of the pointers. Additionally, assigning the return values of the functions to const-qualified pointers will result in the compiler issuing a diagnostic if an attempt is made to modify an object.

Note: If a modified version is required, a program should make and modify a copy of any value covered by this rule.

Example

The following examples are non-compliant as the returned pointers are assigned to non-const qualified pointers. Whilst this will not be reported by a compiler (it is not a constraint violation), an analysis tool will be able to report a violation.

```
void f1 ( void )
{
    char          *s1   = setlocale ( LC_ALL, 0 ); /* Non-compliant */
    struct lconv *conv = localeconv ();          /* Non-compliant */

    s1[ 1 ]       = 'A'; /* Undefined behaviour */
    conv->decimal_point = "^"; /* Undefined behaviour */
}
```

The following examples are compliant as the returned pointers are assigned to const qualified pointers. Any attempt to modify an object through a pointer will be reported by a compiler or analysis tool as this is a constraint violation.

```
void f2 ( void )
{
    char str[ 128 ];

    ( void ) strcpy ( str,
                     setlocale ( LC_ALL, 0 ) ); /* Compliant - 2nd parameter to
                                                    strcpy takes a const char * */
    const struct lconv *conv = localeconv (); /* Compliant */

    conv->decimal_point = "^"; /* Constraint violation */
}
```

The following example shows that whilst the use of a const-qualified pointer gives compile time protection of the value returned by *localeconv*, the same is not true for the strings it references. Modification of these strings can be detected by an analysis tool.

```
void f3 ( void )
{
    const struct lconv *conv = localeconv (); /* Compliant */

    conv->grouping[ 2 ] = 'x'; /* Non-compliant */
}
```

See also

Rule 7.4, Rule 11.8

Rule 21.20 The pointer returned by the Standard Library functions *asctime*, *ctime*, *gmtime*, *localtime*, *localeconv*, *getenv*, *setlocale* or *strerror* shall not be used following a subsequent call to the same function

Category	Mandatory
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

For the purposes of this rule:

- a call to `setlocale` function following a call to `localeconv` function shall be treated as if they are calls to the same function.
- the `asctime` and `ctime` functions shall be treated as if they are the same function.
- the `gmtime` and `localtime` functions shall be treated as if they are the same function.

Rationale

The Standard Library functions *asctime*, *ctime*, *gmtime*, *localtime*, *localeconv*, *getenv*, *setlocale* and *strerror* return a pointer to an object within the Standard Library. Implementations are permitted to use static buffers for any of these objects and a second call to the same function may modify the contents of the buffer. The value accessed through a pointer held by the program before a subsequent call to a function may therefore change unexpectedly.

Example

```
void f1( void )
{
    const char *res1;
    const char *res2;
    char copy[ 128 ];

    res1 = setlocale ( LC_ALL, 0 );

    ( void ) strcpy ( copy, res1 );

    res2 = setlocale ( LC_MONETARY, "French" );

    printf ( "%s\n", res1 ); /* Non-compliant - use after subsequent call */
    printf ( "%s\n", copy ); /* Compliant - copy made before subsequent call */
    printf ( "%s\n", res2 ); /* Compliant - no subsequent call before use */
}
```

Rule 21.21 The Standard Library function *system* of `<stdlib.h>` shall not be used

C90 [Implementation 73]

C99 [Undefined 175; Implementation J.3.2(11), J.3.12(38)]

C11 [Undefined 186; Implementation J.3.2(12), J.3.12(40)]

Category Required

Analysis Decidable, Single Translation Unit

Applies to C90, C99, C11

Amplification

The identifier *system* shall not be used and no macro with this name shall be expanded.

Rationale

This function has undefined and implementation-defined behaviour associated with it.

Errors related to the use of *system* are a common cause of security vulnerabilities.

Rule 21.22 All operand arguments to any *type-generic macros* declared in `<tgmath.h>` shall have an appropriate *essential type*

C99 [Undefined 184], C11 [Undefined 195]

Category Mandatory

Analysis Decidable, Single Translation Unit

Applies to C99, C11

Amplification

The operand arguments passed to the *type-generic* macros defined in `<tgmath.h>` shall have *essentially signed*, *essentially unsigned* or *essentially floating* (either *essentially real floating* or *essentially complex floating*) type.

Arguments to the following macros shall not have *essentially complex floating* type:

atan2, *cbrt*, *ceil*, *copysign*, *erf*, *erfc*, *exp2*, *expm1*, *fdim*, *floor*, *fma*, *fmax*, *fmin*, *fmod*, *frexp*, *hypot*, *ilogb*, *ldexp*, *lgamma*, *llrint*, *llround*, *log10*, *log1p*, *log2*, *logb*, *lrint*, *lround*, *nearbyint*, *nextafter*, *nexttoward*, *remainder*, *remquo*, *rint*, *round*, *scalbn*, *scalbln*, *tgamma*, *trunc*.

Note: The final parameter to the *frexp* and *remquo* macros is for output, and is not considered an operand.

Rationale

Arguments of non-arithmetic types are not convertible to any of the *corresponding real types* defined for the macros defined in `<tgmath.h>`. Attempting to use them therefore results in undefined behaviour.

Casting an argument with *essentially signed* or *essentially unsigned* type to an *essentially real floating* type is not required because the purpose of these macros is to be type-generic. The essential type of the parameter derives from the argument.

Passing an *essentially complex floating* argument to one of the macros listed in the amplification results in undefined behaviour.

Example

Real-valued arguments must have *essentially signed* or *essentially unsigned* or *essentially floating* type:

```
float  f1, f2;
int    i1, i2;

char   c1, c2;
void   *p1, *p2;

void fn1 (void)
{
    f2 = sqrt (f1); /* Compliant      - essentially floating real type */
    i2 = sqrt (i1); /* Compliant      - essentially integer real type */

    c2 = sqrt (c1); /* Non-compliant - essentially character real type */

    p2 = sqrt (p1); /* Non-compliant - undefined behaviour          */
}
```

Arguments to the macros listed in the amplification must have a real type:

```
float f1, f2;
_Complex float cf1, cf2;

void fn2 (void)
{
    f2 = sqrt ( f1); /* Compliant      - real argument          */
    cf2 = sqrt (cf1); /* Compliant      - sqrt has a complex equivalent */

    f2 = ceil ( f1); /* Compliant      - real argument          */
    cf2 = ceil (cf1); /* Non-compliant - undefined behaviour          */
}
```

See also

Rule 21.11, Rule 21.23

Rule 21.23 All operand arguments to any multi-argument *type-generic macros* declared in `<tgmath.h>` shall have the same standard type

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C99, C11

Amplification

This rule applies to the following multi-argument *type-generic macros*:

atan2, copysign, fdim, fma, fmax, fmin, fmod, frexp, hypot, ldexp, nextafter, nexttoward, pow, remainder, remquo, scalbn, scalbln

All operand arguments passed to any of the multi-argument macros defined in `<tgmath.h>` shall have the same standard type, after integer promotion has been applied to any integer arguments.

Note: The final parameter to the *frexp* and *remquo* macros is for output, and is not considered an operand.

Rationale

Ensuring that the types used to determine the *corresponding real type* for the whole call expression are consistent makes the relationship between the input values and the result clearer.

On platforms with extended real types, the *essentially real floating* types may not be able to be strictly ordered by precision. In this case there is a risk that if two arguments have different types, the deduced common real type may be an unexpected size or silently lose precision.

Example

```
void f (void) {
    float32_t f1;
    float32_t f2;
    float64_t d1;
    float64_t d2;

    f2 = pow(f1, f2);          /* Compliant                               */
    d2 = pow(d1, d2);          /* Compliant                               */

    f2 = pow(f1, d2);          /* Non-compliant - unclear which argument was
                                intended to control precision */
    f2 = pow(f1, (float32_t)d2); /* Compliant                               */
}

void g (void) {
    short s16;
    int i32;
    long l32;

    i32 = pow( s16, i32 );      /* Compliant - both arguments are int after
                                integer promotion */
    i32 = pow( i32, l32 );      /* Non-compliant - arguments are not the same
                                type after promotion */

    i32 = pow( s16, 10 );       /* Compliant - 10 has literal type int */
    i32 = pow( 10u, 110ul );    /* Non-compliant - literal types unsigned int
                                and unsigned long */
}

```

See also

Rule 21.11, Rule 21.22

Rule 21.24 The random number generator functions of `<stdlib.h>` shall not be used

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C90, C99, C11

Amplification

The functions `rand` and `srand` shall not be used and no macro with one of these names shall be expanded.

Rationale

The C Standard Library function `rand()` makes no guarantees as to the quality of the random sequence produced.

The C Standard warns that the numbers generated by some implementations of the `rand()` function have a short cycle and the results can be predictable. Applications with particular requirements should use a generator that is known to be sufficient for their needs.

The `srand()` function is also included within this rule, as without any associated use of the `rand()` function, its use is superfluous.

Example

```
#include <stdlib.h>

int r = rand();    /* Non-compliant */
```

Rule 21.25 All memory synchronization operations shall be executed in sequentially consistent order

C11 [Undefined *]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C11

Amplification

The Standard provides an enumerated type `memory_order` to specify the behaviour of memory synchronization operations. Only the memory order `memory_order_seq_cst` shall be used.

The following library functions implicitly use memory ordering *memory_order_seq_cst*:

```
atomic_store, atomic_load, atomic_flag_test_and_set, atomic_flag_clear,
atomic_exchange, atomic_compare_exchange_strong, atomic_compare_exchange_weak,
atomic_fetch_add, atomic_fetch_sub, atomic_fetch_or, atomic_fetch_xor,
atomic_fetch_and
```

For each of these functions, there exists an alternate version with the function name ending in *_explicit()*, which takes an explicit *memory_order* parameter. The functions ending in *_explicit()* shall only be called with the enumeration *memory_order_seq_cst* as the *memory_order* parameter.

Also the following functions shall only be called with the enumeration *memory_order_seq_cst* as the *memory_order* parameter:

```
atomic_thread_fence, atomic_signal_fence
```

Rationale

The Standard defines *memory_order_seq_cst* as the default memory order for objects with atomic types. This ordering is fully defined in the C Standard and enables sequential consistency. The behaviour of other memory orders is non-portable, as it depends on hardware architecture and compiler.

For *memory_order_relaxed*, no operation orders memory. Usage of *memory_order_relaxed* can cause unintuitive behaviour and is error-prone.

Many of those library functions listed above impose restrictions on the memory order allowed, e.g. it is undefined behaviour if the *atomic_store* generic function is called with a *memory_order_acquire*, *memory_order_consume*, or *memory_order_acq_rel* order argument. In case of non-compliant usage, compilers may show warnings but still generate code.

Example

```
typedef struct s {
    uint8_t a;
    uint8_t b;
} s_t;
_Atomic s_t astr;

void main( void )
{
    s_t lstr = {7, 42};

    atomic_init( &astr, lstr );

    lstr = atomic_load( &astr ); /* Compliant */
    lstr = atomic_load_explicit( &astr, memory_order_relaxed ); /* Non-compliant */

    lstr.b = 43;
    atomic_store_explicit( &astr, lstr, memory_order_release ); /* Non-compliant */
}
```

See also

Dir 4.13

Rule 21.26 The Standard Library function *mtx_timedlock()* shall only be invoked on mutex objects of appropriate mutex type

C11 [Undefined *]

Category	Required
Analysis	Undecidable, System
Applies to	C11

Amplification

The first argument of the Standard Library function *mtx_timedlock()* shall be a mutex object of mutex type `mtx_timed` or `(mtx_timed | mtx_recursive)`.

Rationale

Calling the function *mtx_timedlock()* on a mutex object that does not support timeout is undefined behaviour.

Example

```

mtx_t Ra;
mtx_t Rb;
mtx_t Rc;
struct timespec *ts;

void main( void )
{
    mtx_init( &Ra, mtx_plain );
    mtx_init( &Rb, mtx_timed );
    mtx_init( &Rc, mtx_timed | mtx_recursive );
    ...
}

int32_t t1( void* ignore )
{
    ...
    mtx_timedlock( &Ra, ts ); /* Non-compliant */
    mtx_timedlock( &Rb, ts ); /* Compliant */
    mtx_timedlock( &Rc, ts ); /* Compliant */
    ...
}

```

8.22 Resources

Many of the rules in this section are applicable only when rules in other sections have been deviated.

Rule 22.1 All resources obtained dynamically by means of Standard Library functions shall be explicitly released

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

This rule applies to *malloc*, *calloc*, *realloc*, *aligned_alloc* and *fopen*.

Rationale

If resources are not explicitly released then it is possible for a failure to occur due to exhaustion of those resources. Releasing resources as soon as possible reduces the possibility that exhaustion will occur.

Example

```
#include <stdlib.h>

int main ( void )
{
    void *b = malloc ( 40 );

    /* Non-compliant - dynamic memory not released */
    return 1;
}

#include <stdio.h>

int main ( void )
{
    FILE *fp = fopen ( "tmp", "r" );

    /* Non-compliant - file not closed */
    return 1;
}
```

In the following non-compliant example, the handle on "tmp-1" is lost when "tmp-2" is opened.

```
#include <stdio.h>

int main ( void )
{
    FILE *fp;

    fp = fopen ( "tmp-1", "w" );

    fprintf ( fp, "*" );
}
```

```

/* File "tmp-1" should be closed here, but stream 'leaks' */
fp = fopen ( "tmp-2", "w" );
fprintf ( fp, "!" );
fclose ( fp );
return ( 0 );
}

```

See also

Dir 4.12, Dir 4.13, Rule 21.3, Rule 21.6

Rule 22.2 A block of memory shall only be freed if it was allocated by means of a Standard Library function

C90 [Undefined 92], C99 [Undefined 169], C11 [Undefined 179]

Category	Mandatory
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

The Standard Library functions that allocate memory are *malloc*, *calloc* and *realloc*.

A block of memory is freed when its address is passed to *free* and potentially freed when its address is passed to *realloc*. Once freed, a block of memory is no longer considered to be allocated and therefore cannot subsequently be freed again.

Rationale

Freeing non-allocated memory, or freeing the same allocated memory more than once leads to undefined behaviour.

Example

```

#include <stdlib.h>

void fn ( void )
{
    int32_t a;

    /* Non-compliant - a does not point to allocated storage */
    free ( &a );
}

```

```

void g ( void )
{
  char *p = ( char * ) malloc ( 512 );
  char *q = p;

  free ( p );

  /* Non-compliant - allocated block freed a second time */
  free ( q );

  /* Non-compliant - allocated block may be freed a third time */
  p = ( char * ) realloc ( p, 1024 );
}

```

See also

Dir 4.12, Dir 4.13, Rule 21.3

Rule 22.3 The same file shall not be open for read and write access at the same time on different streams

C90 [Implementation 61], C99 [Implementation J.3.12(22)], C11 [Implementation J.3.12(24)]

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

This rule applies to files opened with the Standard Library functions. It may also apply to similar features provided by the execution environment.

Rationale

The C Standard does not specify the behaviour if a file is both written and read via different streams.

Note: it is acceptable to open a file multiple times for read-only access.

Example

```

#include <stdio.h>

void fn ( void )
{
  FILE *fw = fopen ( "tmp", "r+" ); /* "r+" opens for read/write */
  FILE *fr = fopen ( "tmp", "r" ); /* Non-compliant */
}

```

See also

Rule 21.6

Rule 22.4 There shall be no attempt to write to a stream which has been opened as read-only

Category	Mandatory
Analysis	Undecidable, System
Applies to	C90, C99, C11

Rationale

The C Standard does not specify the behaviour if an attempt is made to write to a read-only stream. For this reason it is considered unsafe to write to a read-only stream.

Example

```
#include <stdio.h>

void fn ( void )
{
    FILE *fp = fopen ( "tmp", "r" );

    ( void ) fprintf ( fp, "What happens now?" ); /* Non-compliant */

    ( void ) fclose ( fp );
}
```

See also

Rule 21.6

Rule 22.5 A pointer to a FILE object shall not be dereferenced

Category	Mandatory
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

A pointer to a FILE object shall not be dereferenced directly or indirectly (e.g. by a call to `memcpy` or `memcmp`).

Rationale

Within the section on “files”, the C Standard states that the address of a FILE object used to control a stream may be significant and a copy of the object may not give the same behaviour. This rule ensures that such a copy cannot be made.

The direct manipulation of a FILE object is prohibited as this may be incompatible with its use as a stream designator.

Example

```
#include <stdio.h>

FILE *pf1;
FILE *pf2;
FILE f3;
```

```
pf2 = pf1;      /* Compliant */
f3 = *pf2;     /* Non-compliant */
```

The following example assumes that `FILE *` specifies a complete type with a member named `pos`:

```
pf1->pos = 0;   /* Non-compliant */
```

See also

Rule 21.6

Rule 22.6 The value of a pointer to a `FILE` shall not be used after the associated stream has been closed

C99 [Undefined 140], C11 [Undefined 148]

Category Mandatory

Analysis Undecidable, System

Applies to C90, C99, C11

Rationale

The C Standard states that the value of a `FILE` pointer is indeterminate after a close operation on a stream.

Example

```
#include <stdio.h>

void fn ( void )
{
    FILE *fp;
    void *p;

    fp = fopen ( "tmp", "w" );

    if ( fp == NULL )
    {
        error_action ( );
    }

    fclose ( fp );

    fprintf ( fp, "?" ); /* Non-compliant */
    p = fp;             /* Non-compliant */
}
```

See also

Dir 4.13, Rule 21.6

Rule 22.7 The macro `EOF` shall only be compared with the unmodified return value from any Standard Library function capable of returning `EOF`

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

The value returned by any of these functions shall not be subject to any type conversion if it is later compared with the macro `EOF`. *Note:* indirect type conversions, such as those resulting from pointer type conversions, are included within the scope of this rule.

Rationale

An `EOF` return value from these functions is used to indicate that a stream is either at end-of-file or that a read or write error has occurred. The `EOF` value may become indistinguishable from a valid character code if the value returned is converted to another type. In such cases, testing the converted value against `EOF` will not reliably identify if the end of the file has been reached or if an error has occurred.

If these conditions are to be identified by comparison with `EOF`, the comparison shall be made before any conversion of the value occurs. Alternatively, the Standard Library functions *feof* and *ferror* may be used to directly check the status of the stream, either before or after the conversion takes place.

Example

```
void f1 ( void )
{
    char ch;

    ch = ( char ) getchar ();

    /*
     * The following test is non-compliant. It will not be reliable as the
     * return value is cast to a narrower type before checking for EOF.
     */
    if ( EOF != ( int32_t ) ch )
    {
    }
}
```

The following compliant example shows how *feof()* can be used to check for `EOF` when the return value from *getchar()* has been subjected to type conversion:

```
void f2 ( void )
{
    char ch;

    ch = ( char ) getchar ();

    if ( !feof ( stdin ) )
    {
    }
}
```

```

void f3 ( void )
{
    int32_t i_ch;

    i_ch = getchar ();

    /*
     * The following test is compliant. It will be reliable as the
     * unconverted return value is used when checking for EOF.
     */
    if ( EOF != i_ch )
    {
        char ch;

        ch = ( char ) i_ch;
    }
}

```

Rule 22.8 The value of `errno` shall be set to zero prior to a call to an *errno-setting-function*

Category Required

Analysis Undecidable, System

Applies to C90, C99, C11

Amplification

An *errno-setting-function* is one of the following:

```

ftell, fgetpos, fsetpos, fgetwc, fputwc
strtoimax, strtoumax, strtol, strtoul, strtoll, strtoull, strtod, strtold
wcstoimax, wcstoumax, wcstol, wcstoul, wcstoll, wcstoull, wcstof, wcstod, wcstold
wcrtoimb, wcsrtombs, mbrtowc

```

Any other function which returns error information using `errno` is also an *errno-setting-function*. *Note:* this may include additional functions from the Standard Library, as permitted by the C Standard.

“Prior” requires that `errno` shall be set to zero in the same function and on all paths leading to a call of an *errno-setting-function*. Furthermore, there shall be no calls to functions that may set `errno` in these paths. This includes calls to any function within the Standard Library as these are permitted (but not required) to set `errno`.

Rationale

An *errno-setting-function* writes a non-zero value to `errno` if an error is detected, leaving the value unmodified otherwise. The C Standard includes non-normative advice that “a program that uses `errno` for error checking should set it to zero before a library function call, then inspect it before a subsequent library function call”.

In order that errors can be detected, this rule requires that `errno` shall be set to zero before an *errno-setting-function* is called. Rule 22.9 then requires that `errno` be tested after the call.

Exception

The value of `errno` need not be set to zero when it can be proven to be zero.

Example

```
void f ( void )
{
    errnoSettingFunction1();    /* Non-compliant */

    if ( 0 == errno )
    {
        errnoSettingFunction2(); /* Compliant by exception */

        if ( 0 == errno )
        {
        }
    }

    else
    {
        errno = 0;

        errnoSettingFunction3(); /* Compliant */

        if ( 0 == errno )
        {
        }
    }
}
```

See also

Rule 22.9, Rule 22.10

Rule 22.9 The value of `errno` shall be tested against zero after calling an *errno-setting-function*

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

An *errno-setting-function* is one of those described in Rule 22.8.

The test of `errno` shall occur in the same function on all paths from the call of interest, and before any subsequent function calls.

The results of an *errno-setting-function* shall not be used prior to the testing of `errno`.

Rationale

An *errno-setting-function* writes a non-zero value to `errno` if an error is detected, leaving the value unmodified otherwise. The C Standard includes non-normative advice that “a program that uses `errno` for error checking should set it to zero before a library function call, then inspect it before a subsequent library function call”.

As the value returned by an *errno-setting-function* is unlikely to be correct when `errno` is non-zero, the program shall test `errno` to ensure that it is appropriate to use the returned value.

Exception

The value of `errno` does not have to be tested when the return value of an *errno-setting-function* can be used to determine if an error has occurred.

Example

```
void f1 ( void )
{
    errno = 0;

    errnoSettingFunction1 ();

    someFunction ();          /* Non-compliant - function call */

    if ( 0 != errno )
    {
    }

    errno = 0;

    errnoSettingFunction2 ();

    if ( 0 != errno )        /* Compliant */
    {
    }
}

void f2 ( FILE *f, fpos_t *pos )
{
    errno = 0;

    if ( fsetpos ( f, pos ) == 0 )
    {
        /* Compliant by exception - no need to test errno as no out-of-band error
        reported. */
    }
    else
    {
        /* Something went wrong - errno holds an implementation-defined positive value.
        */
        handleError ( errno );
    }
}
```

See also

Rule 22.8, Rule 22.10

Rule 22.10 The value of `errno` shall only be tested when the last function to be called was an *errno-setting-function*

Category	Required
Analysis	Undecidable, System
Applies to	C90, C99, C11

Amplification

An *errno-setting-function* is one of those described in Rule 22.8.

Rationale

The *errno-setting-functions* are the only functions which are required to set `errno` when an error is detected. Other functions, including those defined in the Standard Library that are not *errno-setting-functions*, may or may not set `errno` to indicate that an error has occurred. The use of `errno` to detect errors within these functions will fail for an implementation that does not set `errno` as it will be left unmodified.

Given that a zero value for `errno` does not therefore guarantee the absence of an error within a function that is not an *errno-setting-function*, its value shall not be tested as the outcome must be considered unreliable.

Example

In the following example:

- `atof` may or may not set `errno` when an error is detected;
- `strtod` is an *errno-setting-function*.

```
void f ( void )
{
    float64_t f64;

    errno = 0;

    f64 = atof ( "A.12" );

    if ( 0 == errno )                /* Non-compliant */
    {
        /* f64 may not have a valid value in here.      */
    }

    errno = 0;

    f64 = strtod ( "A.12", NULL );

    if ( 0 == errno )                /* Compliant      */
    {
        /* f64 will have a valid value in here.        */
    }
}
```

See also

Rule 22.8, Rule 22.9

Rule 22.11 A thread that was previously either joined or detached shall not be subsequently joined nor detached

C11 [Undefined *]

Category	Required
Analysis	Undecidable, System
Applies to	C11

Rationale

Invoking `thrd_detach()` or `thrd_join()` on a thread that has been previously detached or joined is undefined behaviour.

Example

```
void main( void )
{
    thrd_t id1, id2, id3, id4;

    thrd_create( &id1, t1, NULL );
    thrd_create( &id2, t2, NULL );
    thrd_create( &id3, t3, NULL );
    thrd_create( &id4, t4, NULL );

    thrd_join ( id1, NULL ); /* Compliant */
    thrd_join ( id1, NULL ); /* Non-compliant - already joined */

    thrd_detach( id2 ); /* Compliant */
    thrd_detach( id2 ); /* Non-compliant - already detached */

    thrd_join ( id3, NULL ); /* Compliant */
    thrd_detach( id3 ); /* Non-compliant - already joined */

    thrd_detach( id4 ); /* Compliant */
    thrd_join ( id4, NULL ); /* Non-compliant - already detached */
}
```

Rule 22.12 Thread objects, thread synchronization objects, and thread-specific storage pointers shall only be accessed by the appropriate Standard Library functions

C11 [Undefined *]

Category Mandatory

Analysis Undecidable, System

Applies to C11

Amplification

Thread objects shall exclusively be accessed via the Standard Library functions *thrd_create()*, *thrd_detach()*, *thrd_join()*, and *thrd_equal()*.

Mutex objects shall exclusively be accessed via the Standard Library functions *mtx_destroy()*, *mtx_init()*, *mtx_lock()*, *mtx_trylock()*, *mtx_timedlock()*, *mtx_unlock()*, *cnd_wait()*, and *cnd_timedwait()*.

Condition variables shall exclusively be accessed via the Standard Library functions *cnd_broadcast()*, *cnd_destroy()*, *cnd_init()*, *cnd_signal()*, *cnd_wait()*, and *cnd_timedwait()*.

Thread-specific storage pointers shall exclusively be accessed by the Standard Library functions *tss_create()*, *tss_delete()*, *tss_get()*, and *tss_set()*.

Rationale

Thread objects and thread synchronization objects are expected to be unique for the corresponding thread and synchronization resources.

Thread-specific storage pointers are identified by unique keys. Any direct manipulation (copy, assignment, etc.) may result in undefined behaviour. The *tss_delete()*, *tss_get()* and *tss_set()* functions shall only be called with a value for key that was returned by a call to *tss_create()*, otherwise the behaviour is undefined.

Example

```

mtx_t  Ra;
mtx_t  Rb;
thrd_t id1;
thrd_t id2;
tss_t  key;

int32_t t1( void *ignore )
{
    mtx_lock( &Ra );
    int32_t val;
    if ( id1 == id2 )          /* Non-compliant - use thrd_equal()      */
    {
        Rb = Ra;              /* Non-compliant                    */
        memcpy(&Rb, &Ra, sizeof(mtx_t)); /* Non-compliant                    */
    }

    if ( thrd_equal( id1, id2 ) ) /* Compliant                        */
    {
        ...
    }
    key++;                    /* Non-compliant - explicit manipulation of
                               TSS pointer                                */
    tss_set( key, &val );     /* Undefined, value of key not returned by
                               tss_create()                                */
}

void main( void )
{
    mtx_init ( &Ra, mtx_plain );
    mtx_init ( &Rb, mtx_plain );
    tss_create ( &key, NULL );
    thrd_create( &id1, t1, NULL );
    thrd_create( &id2, t1, NULL );
    ...
}

```

See also

Rule 11.5, Rule 22.20

Rule 22.13 Thread objects, thread synchronization objects and thread-specific storage pointers shall have appropriate storage duration

C11 [Undefined 9, 10, 11]

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C11

Amplification

Objects of type *thrd_t*, *mtx_t*, *cond_t*, and *tss_t* shall not have automatic storage duration nor thread storage duration.

Rationale

Determining the lifetime of non-static objects which depend on thread execution state is difficult and error-prone. In particular, sharing objects of automatic storage duration between threads and using them to control concurrent execution can cause undefined behaviour due to accessing them outside of their lifetime.

Usage of a static pool of synchronization resources is common practice in many safety-related operating systems, including ARINC-653 [45], AUTOSAR [46] and OSEK [47].

Example

```

mtx_t Ra;                                     /* Compliant */

int32_t t1( void *ptr )                       /* Thread entry */
{
    ...
    mtx_lock ( &Ra );
    mtx_lock ( (mtx_t*)ptr );                 /* Lifetime of Rb might have ended
                                              ... pointer might be dangling */
    ...
    mtx_unlock( (mtx_t*)ptr );               /* Lifetime of Rb might have ended
                                              ... pointer might be dangling */
    mtx_unlock( &Ra );
}

void main( void )
{
    thrd_t id1;                               /* Non-compliant */
    mtx_t Rb;                                 /* Non-compliant */

    mtx_init ( &Ra, mtx_plain );
    mtx_init ( &Rb, mtx_plain );
    thrd_create( &id1, t1, &Rb );
}

```

Rule 22.14 Thread synchronization objects shall be initialized before being accessed

C11 [Undefined 9]

Category Mandatory

Analysis Undecidable, System

Applies to C11

Amplification

Before being accessed, mutex objects shall be initialized by calling *mtx_init()*, and condition variables by calling *cnd_init()*.

The second parameter of *mtx_init()* shall be either *mtx_plain*, *mtx_timed*, (*mtx_plain* | *mtx_recursive*), or (*mtx_timed* | *mtx_recursive*).

Rationale

Mutex objects have to be explicitly created by calling function *mtx_init()*, and condition variables have to be explicitly created by calling function *cnd_init()*.

Invoking `mtx_init()` with a different value of its type parameter than listed above is undefined behaviour.

Initializing all synchronization objects before creating the threads accessing them is a deterministic way to prevent threads from accessing synchronization objects with indeterminate state.

Example

```

mtx_t Ra;
mtx_t Rb;
mtx_t Rc;

int32_t t1( void *ignore )      /* Thread T1 entry */
{
    mtx_init( &Rb, mtx_plain ); /* Non-compliant - T2 may have already accessed Rb */
    ...
    /* Subsequently locks/unlocks Ra, Rb, Rc */
}

int32_t t2( void *ignore )
{
    /* locks/unlocks Ra, Rb, Rc */
}

thrd_t id1, id2;

void main(void)
{
    mtx_init ( &Ra, mtx_plain ); /* Compliant */

    thrd_create( &id1, t1, NULL );
    thrd_create( &id2, t2, NULL );

    mtx_init ( &Rc, mtx_plain ); /* Non-compliant - T1/T2 may have already
                                   accessed Rc */

    thrd_join ( id1, NULL );
    thrd_join ( id2, NULL );

    mtx_destroy( &Ra );
    mtx_destroy( &Rb );
    mtx_destroy( &Rc );
}

```

See also

Dir 4.7

Rule 22.15 Thread synchronization objects and thread-specific storage pointers shall not be destroyed until after all threads accessing them have terminated

C11 [Undefined 9, 10, *]

Category	Required
Analysis	Undecidable, System
Applies to	C11

Rationale

The Standard Library function *mtx_destroy(mtx)* releases all resources used by the mutex pointed to by *mtx*. Destroying a mutex which is still locked by some thread results in undefined behaviour, as the C Standard expects no threads to be blocked by a mutex when it is destroyed.

The Standard Library function *tss_delete(key)* releases all resources used by the thread-specific storage identified by *key*. Calling the *tss_delete()*, *tss_get()* or *tss_set()* functions after the thread commenced executing destructors results in undefined behaviour.

Calling the Standard Library function *cond_destroy()*, on a condition variable on which a thread is currently waiting, results in undefined behaviour.

These problems are avoided by only destroying synchronization resources and deleting thread-specific storage after all threads accessing them have terminated (or not at all).

Example

```

mtx_t  Ra;
mtx_t  Rb;
tss_t  key1;
tss_t  key2;
thrd_t  id1;
thrd_t  id2;

int32_t t1( void *ignore )  /* Thread T1 entry */
{
    /*
    ** locks/unlocks Ra, Rb
    ** accesses thread-specific storage pointed to by key1, key2
    */

    tss_delete( key1 );      /* Non-compliant - might still be accessed from T2 */
}

int32_t t2( void *ignore )  /* Thread T2 entry */
{
    /*
    ** locks/unlocks Ra, Rb
    ** accesses thread-specific storage pointed to by key1, key2
    */

    mtx_destroy( &Rb );     /* Non-compliant - T1 might still access Rb */
}

```

```

void main( void )
{
    mtx_init    ( &Ra, mtx_plain );
    mtx_init    ( &Rb, mtx_plain );

    tss_create ( &key1, NULL );
    tss_create ( &key2, NULL );

    thrd_create( &id1, t1, NULL );
    thrd_create( &id2, t2, NULL );

    spendSomeTime();

    tss_delete ( key2 );          /* Non-compliant - might still be accessed by t1, t2 */

    thrd_join  ( id1, NULL );
    thrd_join  ( id2, NULL );

    mtx_destroy( &Ra );          /* Compliant */
    tss_delete ( key1 );          /* Compliant */
}

```

See also

Rule 22.1

Rule 22.16 All mutex objects locked by a thread shall be explicitly unlocked by the same thread

C11 [Undefined *]

Category	Required
Analysis	Undecidable, System
Applies to	C11

Amplification

If a mutex object *mtx* is locked by *mtx_lock(mtx)* at a program point *p* there shall be an explicit *mtx_unlock(mtx)* for mutex object *mtx* on all programs paths reachable from *p* before exiting the thread.

Rationale

When a thread terminates without releasing a lock, that lock may be held for indeterminate time. If the life range of a mutex object ends while there are threads waiting for it the behaviour is undefined.

Destroying a mutex on which threads are waiting is undefined behaviour.

Note: it is good practice to unlock mutexes in the same function and under the same control dependences in which they have been locked.

Example

```

mtx_t Ra;
mtx_t Rb;

```

```
int32_t t1( void *ignore ) /* Thread 1 */
{
    bool_t b;

    mtx_lock ( &Ra ); /* Compliant */
    mtx_unlock( &Ra );

    mtx_lock ( &Rb ); /* Non-compliant - unlock missing on one path */
    if ( b )
    {
        mtx_unlock( &Rb );
    }
    return 0;
}
```

See also

Dir 4.13, Rule 22.1

Rule 22.17 No thread shall unlock a mutex or call *cond_wait()* or *cond_timedwait()* for a mutex it has not locked before

C11 [Undefined *]

Category	Required
Analysis	Undecidable, System
Applies to	C11

Amplification

A mutex shall only be unlocked by a thread if it has been locked by that thread before.

The *cond_wait()* and *cond_timedwait()* functions shall only be called by a thread on a mutex that is locked by that thread.

Rationale

Unlocking a mutex which has not been locked by the calling thread is undefined behaviour. Calling *cond_wait()* or *cond_timedwait()* with mutex argument *mtx* requires that the mutex pointed to by *mtx* be locked by the calling thread.

Example

```
mtx_t Ra;
mtx_t Rb;
cond_t Cnd1;
cond_t Cnd2;
```

```

int32_t t1( void *ignore ) /* Thread 1 */
{
    mtx_lock ( &Ra );
    mtx_unlock( &Ra ); /* Compliant */

    mtx_unlock( &Ra ); /* Non-compliant - mutex is not locked */

    cnd_wait ( &Cnd1, &Ra ); /* Non-compliant - mutex is not locked */

    mtx_unlock( &Rb); /* Non-compliant - mutex either not locked, or
                       ... is locked by different thread */

    cnd_wait ( &Cnd2, &Rb ); /* Non-compliant - mutex either not locked, or
                              ... is locked by different thread */

    return 0;
}

int32_t t2( void *ignore ) /* Thread 2 */
{
    mtx_lock ( &Rb );
    doSomething();
    mtx_unlock ( &Rb ); /* Compliant */
    return 0;
}

```

See also

Dir 4.13, Rule 22.1, Rule 22.18

Rule 22.18 Non-recursive mutexes shall not be recursively locked

C11 [Undefined *]

Category	Required
Analysis	Undecidable, System
Applies to	C11

Amplification

A non-recursive mutex shall only be locked by a thread if it has not already been locked by that before.

Rationale

It is undefined behaviour if a non-recursive mutex is recursively locked by the calling thread. If the thread also attempts to unlock the mutex twice, the second call to *mtx_unlock()* will also result in undefined behaviour, since the mutex then will already be unlocked.

Example

```

mtx_t Ra;
mtx_t Rb;

```

```

int32_t t1( void *ignore ) /* Thread 1 */
{
    mtx_lock ( &Rb ); /* Compliant */
    mtx_lock ( &Rb ); /* Compliant - Rb is recursive */
    mtx_unlock( &Rb ); /* Rb still locked */
    mtx_unlock( &Rb ); /* Rb gets unlocked */

    mtx_lock ( &Ra ); /* Compliant */
    mtx_lock ( &Ra ); /* Non-compliant - undefined behaviour, deadlock possible */
    mtx_unlock( &Ra ); /* If reachable (i.e. no deadlock), Ra gets unlocked */
    mtx_unlock( &Ra ); /* Undefined behaviour if reachable */

    return 0;
}

thrd_t id1;
thrd_t id2;

int32_t main(void)
{
    mtx_init ( &Ra, mtx_plain );
    mtx_init ( &Rb, mtx_recursive );
    thrd_create( &id1, t1, NULL );
    ...
}

```

See also

Dir 4.13, Rule 22.1, Rule 22.17

Rule 22.19 A condition variable shall be associated with at most one mutex object

C11 [Undefined *]

Category	Required
Analysis	Undecidable, System
Applies to	C11

Rationale

If the same condition variable is used with different mutex objects by two threads, it is undefined which mutex will be unlocked upon signalling the condition variable.

Example

```

mtx_t Ra;
mtx_t Rb;
cnd_t Cnd;

int32_t t1(void *ignore )
{
    mtx_lock ( &Ra );
    cnd_wait ( &Cnd, &Ra ); /* Non-compliant - t2 uses Cnd with Rb */
    mtx_unlock( &Ra );
    return 0;
}

```

```

int32_t t2(void *ignore )
{
    mtx_lock ( &Rb );
    cnd_wait ( &Cnd, &Rb ); /* Non-compliant - t1 uses Cnd with Ra */
    mtx_unlock( &Rb );
    return 0;
}

int32_t t3(void* ignore)
{
    cnd_signal( &Cnd ); /* Unblocks one of Ra and Rb...
                        ... unclear whether t1 or t2 resumes */
    return 0;
}

```

Rule 22.20 Thread-specific storage pointers shall be created before being accessed

C11 [Undefined 9, *]

Category Mandatory

Analysis Undecidable, System

Applies to C11

Amplification

Objects of type *tss_t* shall be explicitly created by *tss_create()* before being accessed.

Rationale

Thread-specific storage pointers have to be explicitly created before accessing them. Creating them inside of threads creates dependencies on thread execution and ordering which are hard to maintain and check. Creating them before creating the threads accessing them is a deterministic way to prevent threads from accessing thread-specific storage pointers with indeterminate state.

Example

```

tss_t key1;
tss_t key2;
thrd_t id1;
thrd_t id2;
int32_t g1;
int32_t g2;

int32_t t2( void *ignore ) /* Thread t2 entry */
{
    tss_create( &key1, NULL ); /* Non-compliant - thread t1 might already have
                                tried to access key1 */
}

```

```

int32_t t1( void *ignore )      /* Thread t1 entry */
{
    tss_set    ( key1, &g1 );   /* Non-compliant - might not yet be created */
    tss_set    ( key2, &g2 );   /* Compliant */

    int32_t *v1 = tss_get( key1 );
    int32_t *v2 = tss_get( key2 );

    *v1 = computeG1();
    *v2 = computeG2();
}

void main( void )
{
    int32_t i;

    tss_create( &key2, NULL );  /* Compliant */

    thrd_create( &id1, t1, NULL );
    thrd_create( &id2, t2, NULL );
}

```

See also

Dir 4.13

8.23 Generic Selections

Rule 23.1 A generic selection should only be expanded from a macro

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C11

Amplification

A generic selection expression should only be expanded from a function-like macro, with one of the macro's arguments expanded into the controlling expression of the generic selection.

Rationale

Generic selections are useful to implement type queries or checks, or generic functions. If the generic selection is applied directly rather than expanded from a macro body, the type of the operand is already known locally and querying it is not necessary.

Example

The following example is non-compliant, as it is not a macro:

```

int32_t x = 0;

/* Non-compliant - not a macro */
int32_t y = _Generic( x
    , int32_t    : 1
    , float32_t : 2 );

```

The following example implements a regular generic function, and can be expanded into contexts where the operand type may vary:

```
/* Compliant - used to implement a generic function */
#define arith(X) ( _Generic( (X)
                  , int32_t   : handle_int32
                  , float32_t : handle_float
                  , default   : handle_any   ) (X) )
```

The following example is not compliant because it expands to a generic selection which does not depend on the type of an argument:

```
/* Non-compliant */
#define maybe_inc(Y) ( _Generic( x
                               , int32_t : 1
                               , default : 0 ) + (Y) )
```

See also

Dir 4.9, Rule 13.6, Rule 23.2

Rule 23.2 A generic selection that is not expanded from a macro shall not contain potential *side effects* in the controlling expression

Category Required

Analysis Decidable, Single Translation Unit

Applies to C11

Amplification

If the controlling expression of a generic selection is not expanded from a macro argument, it shall not appear to contain any *side effects*.

A function call is considered to be a *side effect* for the purposes of this rule.

Rationale

The controlling expression of a generic selection is never evaluated. It is only used for its type, and usually has to be repeated in order to be combined with the result value in some way.

If an expression is specified which syntactically contains a side effect, that effect will not be applied, even if the expression has a type that would cause it to be evaluated by `sizeof`, such as a VLA.

Example

```
#ifdef BIG
typedef int64_t STATE;
#else
typedef int16_t STATE;
#endif

STATE shared_state;

/* Compliant */
_Static_assert ( _Generic ( shared_state
                          , int32_t: 0
                          , default: 1 )
               , "error on wrong type");
```

This example is non-compliant because the apparent side effect modifying `shared_state` is not evaluated by the controlling expression:

```
/* Non-compliant */
_Static_assert ( _Generic ( ++shared_state
                        , int32_t: 0
                        , default: 1)
                , "error on wrong type");
```

See also

Rule 13.6, Rule 23.1, Rule 23.7

Rule 23.3 A generic selection should contain at least one non-default association

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C11

Amplification

A generic selection should contain at least one association which explicitly specifies an object type. The presence of a `default` association is always optional.

Rationale

A generic selection that consists only of a usable `default` association does not do anything useful; the default association's result value is evaluated unconditionally.

Omitting a default association indicates intent to check the operand type, by introducing a constraint violation when the type does not match.

Example

```
/* Non-compliant - consists only of a default association */
#define no_op(X) _Generic( (X), default: (X) )

/* Compliant - has a non-default and a default association */
#define filter_ints(X) ( _Generic( (X) \
                                , int32_t: handle_int \
                                , default: handle_numeric_value ) (X) )

/* Compliant - has non-default associations */
#define only_ints(X) ( _Generic( (X) \
                                , int32_t: handle_int \
                                , uint32_t: handle_int ) (X) )
```

The following example demonstrates that, by omitting a default association, there is a clear intent to check the operand type, by introducing a constraint violation when the type does not match:

```
/* Compliant - it has a single permitted type and is
   intended to prevent implicit conversion */
#define require_char(X) ( _Generic ( (X), char8_t: (X) ) )
```

See also

Rule 23.4, Rule 23.8

Rule 23.4 A generic association shall list an appropriate type

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C11

Amplification

The controlling expression of a generic selection undergoes *lvalue conversion* before having its type compared to the entries in the association list. The association list shall not contain any associations for:

- A `const`-qualified object type
- A `volatile`-qualified object type
- An atomic object type
- An array type
- A function type
- An unnamed structure or union definition

Rationale

Since the C Standard does not impose a constraint limiting the types in the association list to the possible types of values after conversion, it is possible to list associations for types that would never be eligible for selection.

Value conversion removes top-level qualification and “decays” array and function values into pointers. Therefore, such types can never match the type of the converted value of the controlling expression. Listing them is not a constraint violation, but serves no useful purpose, and is almost certainly an error. This only affects object types, not qualification on pointed-to types.

Listing an unnamed structure or union in the association list is a violation of this rule because every occurrence of a *struct-declaration-list* creates a distinct type, therefore it can only be *matched* by a default association.

Example

In this set of examples, the first pair are non-compliant because they explicitly specify generic associations for types that can never be the type of the controlling expression, while the second pair are compliant as they only attempt to specify generic associations for types that can match the controlling expression:

```
typedef int32_t Func (int);
typedef int32_t Array [10];

typedef Func * FuncP;
typedef Array * ArrayP;

/* Non-compliant (because Func is listed) */
#define handle_function_nc(X) _Generic((X) \
, Func : handle_funcp (&(X)) \
, FuncP: handle_funcp (X) )
```

```

/* Non-compliant (because Array is listed) */
#define handle_array_nc(X) _Generic((X)
    , Array      : handle_intp ((X) + 0) \
    , ArrayP     : handle_intp (*(X))   \
    , int32_t *: handle_intp (X) )

/* Compliant */
#define handle_function(X) _Generic( (X), FuncP: handle_funcp (X) )

/* Compliant */
#define handle_array(X) _Generic(&(X)[0]
    , int32_t *: handle_intp (X) \
    , ArrayP   : handle_intp (*(X)) )

Array arr1;
Array arr2[10]; /* Two-dimensional - array of arrays */

handle_array (arr1); /* type of &(arr1[0]) is int32_t * - can pass to handle_intp */
handle_array (arr2); /* type of &(arr2[0]) is ArrayP - must dereference again */

```

In this set of examples, the first example is non-compliant because it explicitly specifies a type that includes qualifiers at the object level; qualification is removed by value conversion, so such types are not selectable. The second example is able to take qualification into account because the qualifiers apply to the pointed-to type, not the object type, and are not removed from the type of the value of the controlling expression:

```

typedef int32_t Int;
typedef int32_t const CInt;

/* Non-compliant */
#define filter_const_nc(X) ( _Generic((X)
    , CInt      : handle_const_intp \
    , default: handle_other_value ) (&(X)) )

/* Compliant */
#define filter_const(X) ( _Generic((X) \
    , CInt * : handle_const_intp \
    , Int *  : handle_const_intp \
    , default: handle_other_value ) (X) )

```

In this set of examples, the first example is non-compliant because it tries to explicitly specify array types, which are not selectable, in an attempt to involve dependent type information in the match. The second example below is compliant because it explicitly specifies a pointer type, and forces the type of the controlling expression to be a pointer with &. Creating a pointer to an array preserves the complete array type, including its dimension.

```

/* Avoid multiple associations if size == 1 */
#define SizeofNonEmpty(A) (sizeof (A) > 1 ? sizeof (A) : 2)

/* Non-compliant - tries to match argument array type */
#define only_strings_nc(X) _Generic((X)
    , char[SizeofNonEmpty (X)]: handle_sized_string (sizeof (X), (X)) \
    , char[1]                  : handle_null_terminator (1, (X)))

/* Compliant - matches pointer to argument array type */
#define only_strings(X) _Generic(&(X)
    , char (*) [SizeofNonEmpty (X)]: handle_sized_string (sizeof (X), (X)) \
    , char (*) [1]: handle_null_terminator (1, (X)))

only_strings_nc ("hello"); /* Constraint violation */
only_strings    ("world"); /* Matches as char (*) [6] */

```

See also

Rule 23.5

Rule 23.5 A generic selection should not depend on implicit pointer type conversion

Category	Advisory
Analysis	Decidable, Single Translation Unit
Applies to	C11

Amplification

This rule only applies to selection based on pointer types. Selection based on an arithmetic type is never a violation of this rule.

This rule applies when a default association is selected for a controlling expression whose type could, in other contexts (such as when passed as an argument to a function), be implicitly converted to another type explicitly listed in the association list.

Rationale

The controlling expression of a generic selection undergoes value conversion before having its type compared to any of the entries in the association list, but this is the only conversion applied to its type. It can then only match an association exactly, or select the default association.

No attempt is made to implicitly convert the type to match an association.

For instance, listing an association for `void *` will never match a pointer to a complete object type, and listing an association for `T const *` will never match `T *`. This is different from function arguments, and may therefore cause surprise for users of a generic function implemented as a generic association, if a type that would be implicitly convertible to an explicitly-listed type is instead allowed to fall through to the default association.

Example

Given:

```
void handle_pi (int32_t *); /* No associations have this signature */
void handle_cpi (const int32_t *);
void handle_any (void *);
```

Using the following generic function with a pointer to `const int32_t` is compliant, but using a pointer to (mutable) `int32_t` is not, because there is no implicit conversion to qualified-pointer in generic matching:

```
#define handle_pointer1(X) ( _Generic ((X) \
, const int32_t *: handle_cpi \
, default : handle_any) (X) )

const int32_t ci;
handle_pointer1 (&ci); /* Compliant - const int32_t * is selected */

int32_t mi;
handle_pointer1 (&mi); /* Non-compliant - default is selected
                        NOT const int32_t * */
```

Using the following generic function with a (mutable) pointer to `int32_t` is non-compliant, as there is no conversion to qualified pointer or to pointer to `void` in generic matching.

```
#define handle_pointer2(X) ( _Generic ((X)
                                , void      *: handle_any
                                , const int32_t *: handle_cpi
                                , default    : handle_any) (X))

handle_pointer2 (&mi); /* Non-compliant - default is selected
                        NOT void * or const int32_t * */
```

A generic function that wishes to handle pointers to any object type and distinguish them by qualification, can wrap the controlling expression such that the pointed-to type is explicitly converted to a pointer to `void`:

```
void handle_unq (void      *);
void handle_cq  (void const *);
void handle_vq  (void      volatile *);
void handle_cvq (void const volatile *);

#define handle_pointer(X) ( _Generic (1 ? (X) : (void *) (X)
                                , void      * : handle_unq
                                , void const * : handle_cq
                                , void      volatile * : handle_vq
                                , void const volatile * : handle_cvq) (X) )

/* Uses are always compliant */
```

Per the specification of the conditional operator's composite result type, any argument to the macro with a pointer to object type will be converted to a similarly-qualified pointer to `void` before it is used as the controlling expression.

See also

Rule 23.3, Rule 23.4, Rule 23.6

Rule 23.6 The controlling expression of a generic selection shall have an *essential type* that matches its standard type

Category Required

Analysis Decidable, Single Translation Unit

Applies to C11

Amplification

Using an expression that has a distinct essential type from its standard type as the controlling expression of a generic selection is always a violation of this rule, except for integer constant expressions that are neither character constants nor essentially boolean.

Notes:

1. An enumeration constant unconditionally has the type `signed int` and will not match either the implementation defined underlying type, if different from `int`, or the (indistinguishable) enumerated type, whichever is listed.
2. The same consideration applies to `true` and `false` defined in `<stdbool.h>`, which have type `int` and will not match an association of type `_Bool`.

Rationale

The association selected by a generic selection must be in line with the type system under which the code was designed, in order to be useful and consistent. Code written under MISRA guidelines is subject to the essential type system, and therefore if the selected type is not consistent with the essential type of the controlling expression, the generic selection will have violated the type system used for the design.

Because generic selections choose a standard type by name, they can expose this difference if used with an argument that has inconsistent essential and standard types.

Exception

This rule does not apply to integer constant expressions which would have an essentially signed or unsigned type of lower rank than `int`, because of the *Type of Lowest Rank* rule, but are neither character constants nor essentially boolean.

Example

The non-compliant examples below accept an operand with essential type of `signed short` and `char` respectively, but both have a standard type of `signed int`, and only the standard type determines which association will match.

An exception is made for integer constant expressions that are not character constants, because this is in line with most user expectations.

```
#define filter_ints(X) ( _Generic((X)
    , signed short: handle_sshort \
    , unsigned short: handle_ushort \
    , signed int : handle_sint \
    , unsigned int : handle_uint \
    , signed long : handle_slong \
    , unsigned long : handle_ulong \
    , default : handle_any) (X) )

short s16 = 0;
int i32 = 0;
long l32 = 0;

/* Non-compliant */
filter_ints (s16 + s16); /* Selects int, essentially short */
filter_ints ('c'); /* Selects int, essentially char */

/* Compliant */
filter_ints (s16);
filter_ints (i32);
filter_ints (l32);

/* Compliant by exception */
filter_ints (10u); /* UTLR is unsigned char */
filter_ints (250 + 350); /* STLR is signed short */
```

Because an enumerated type is compatible with its implementation-defined underlying type, it is not distinguishable from the underlying type by `_Generic` (although distinguishing different enumerated types is possible because this compatibility is not transitive).

Attempting to list both would be a constraint violation. Therefore, using an enumerated type as the controlling expression to a generic selection that lists the underlying type, or vice versa, is always a violation.

```
enum E { A = 0, B = 1, C = 2 };
enum E e = A;
```

```
/* Non-compliant, assuming compiler chooses 'unsigned int' as the underlying type */
filter_ints (e); /* selects 'unsigned int', essentially enum */
filter_ints (A); /* selects 'signed int', essentially enum */
```

See also

Rule 23.5

Rule 23.7 A generic selection that is expanded from a macro should evaluate its argument only once

Category Advisory

Analysis Decidable, Single Translation Unit

Applies to C11

Amplification

If the controlling expression of a generic selection is expanded from a macro argument, it should also be expanded elsewhere in the macro body so that it is evaluated exactly once in the entire expanded macro body. The number of times that the expression is evaluated should not depend on which association is selected, and should be consistent for all associations in the generic selection.

This rule does not depend on the presence of side effects in the operand expression.

Rationale

The controlling expression of a generic selection is never evaluated. It is only used for its type, and usually has to be repeated in order to be combined with the result value in some way.

If an expression is specified which syntactically contains a side effect, that effect will not be applied. If the generic selection is (as is usually the case) the result of a macro expansion which uses one of the macro's arguments to select the type, this non-evaluation is concealed from the invoking code.

Combining the input value with the result expression “outside” the generic selection (such as choosing a function to call as the result, rather than the complete function call) is more likely to guarantee consistent evaluation of the expanded operand.

Exception

This rule is not violated if all result expressions for the generic selection are constant expressions, and the macro never expands the argument outside of the controlling expression. This allows for the implementation of type queries like `is_pointer_const`.

Example

The following examples are consistent in their expansion of side effects for all associations. In the first example, the argument used for the controlling expression are expanded exactly once outside of the generic selection, so the generic selection does not prevent it from being evaluated. In the second example, the argument expands once each into the result expression of each association. It will be evaluated exactly once regardless of which generic association is *selected*.

```
/* Compliant */
#define gfun1(X) ( _Generic((X)           \
                  , float32_t: fun1f     \
                  , float64_t: fun1      \
                  , default_ : fun1l) (X) )

#define gfun2(X)  _Generic((X)           \
                          , float32_t: fun2f (X) \
                          , float64_t: fun2 (X) \
                          , default_ : fun2l (X) )
```

The following example is non-compliant because whether or not the macro argument is evaluated depends on whether the default association is *selected* or not. The default association does not evaluate it because it has no use for the value, but this dangerously assumes that there were no side effects in the expression.

```
/* Non-compliant */
#define gfun3(X)  _Generic((X)           \
                          , float32_t : fun3f (X) \
                          , float64_t : fun3 (X) \
                          , float128_t: fun3l (X) \
                          , default_  : default_result )
```

The following example only expands a macro argument into the controlling expression, and all possible result expressions are integer constant expressions. This implements a type query that can be used for type checking.

```
/* Compliant by exception */
#define is_pointer_const(P) _Generic(1 ? (P) : (void *) (P) \
                                     , void const          * : 1 \
                                     , void const volatile * : 1 \
                                     , default              : 0 )

_Static_assert (is_pointer_const (pi), "must not be an out-parameter");
```

See also

Rule 23.2

Rule 23.8 A default association shall appear as either the first or the last association of a generic selection

Category	Required
Analysis	Decidable, Single Translation Unit
Applies to	C11

Amplification

A default association shall only appear as either the first generic association in the association list, or the last – this rule only applies when a generic selection contains a default association.

Rationale

Having a default association appear as either the first or last in the association list makes it easy to locate, clarifying the intent of the selection structure.

Example

```

/* Non-compliant - default is not first or last association */
#define sqrt(X) ( _Generic((X)
                  , float32_t : sqrtf
                  , default   : sqrt
                  , float128_t: sqrtl) (X) )

/* Compliant - default is first association */
#define cbrt(X) ( _Generic((X)
                          , default   : cbrt
                          , float32_t : cbrtf
                          , float128_t: cbrtl) (X) )

/* Compliant - no default */
#define assert_untyped_nonatomic(X) ( _Generic((X)
                                               , void
                                               , void const
                                               , void volatile
                                               , void const volatile * : handle_ptr) (X) )

```

See also

Rule 16.5, Rule 23.3

9 References

9.1 MISRA Publications

9.1.1 MISRA C

- [1] MISRA *Guidelines for the use of the C language in vehicle based software*
ISBN 0-9524159-9-0, Motor Industry Research Association, Nuneaton, April 1998
- [2] MISRA *Guidelines for the use of the C language in critical systems*
ISBN 0-9524156-2-3, MIRA Limited, Nuneaton, October 2004
- [3] MISRA *Guidelines for the use of the C language in critical systems* (Third edition)
ISBN 978-1-906400-10-1 paperback, 978-1-906400-11-8 PDF, MIRA Limited, Nuneaton, May 2013
- [4] MISRA *Guidelines for the use of the C language In critical systems* (Third edition, First revision)
ISBN 978-1-906400-21-7 paperback, 978-1-906400-22-4 PDF, HORIBA MIRA Limited, Nuneaton, February 2019
- [5] MISRA C:2012 Amendment 1, *Additional security guidelines for MISRA C:2012*
ISBN 978-1-906400-16-3 PDF, HORIBA MIRA Limited, Nuneaton, April 2016
- [6] MISRA C:2012 Amendment 2, *Updates for ISO/IEC 9899:2011 core functionality*
ISBN 978-1-906400-25-5 PDF, HORIBA MIRA Limited, Nuneaton, February 2020
- [7] MISRA C:2012 Amendment 3, *Updates for ISO/IEC 9899:2011/2018 Phase 2 — New C11/C18 features*
ISBN 978-1-911700-02-9 PDF, The MISRA Consortium Limited, Norwich, October 2022
- [8] MISRA C:2012 Amendment 4, *Updates for ISO/IEC 9899:2011/2018 Phase 3 — Multi-threading and atomics*
ISBN 978-1-911700-03-6 PDF, The MISRA Consortium Limited, Norwich, March 2023
- [9] MISRA C:2012 Technical Corrigendum 1
ISBN 978-1-906400-17-0 PDF, HORIBA MIRA Limited, Nuneaton, June 2017
- [10] MISRA C:2012 Technical Corrigendum 2
ISBN 978-1-911700-00-5 PDF, The MISRA Consortium Limited, Norwich, March 2022

9.1.2 MISRA Compliance

- [11] MISRA Compliance:2020, *Achieving compliance with MISRA coding guidelines*
ISBN 978-1-906400-26-2 PDF, HORIBA MIRA Limited, February 2020

9.1.3 MISRA Auto-Generated Code

- [12] MISRA AC AGC *Guidelines for the application of MISRA-C:2004 in the context of automatic code generation*
ISBN 978-1-906400-02-6, MIRA Limited, Nuneaton, November 2007
- [13] MISRA AC GMG *Generic modelling design and style guidelines*
ISBN 978-1-906400-06-4, MIRA Limited, Nuneaton, May 2009

- [14] MISRA AC SLSF *Modelling design and style guidelines for the application of Simulink and Stateflow* ISBN 978-1-906400-07-1, MIRA Limited, Nuneaton, May 2009
- [15] MISRA AC TL *Modelling style guidelines for the application of Targetlink in the context of automatic code generation* ISBN 978-1-906400-01-9, MIRA Limited, Nuneaton, November 2007

9.1.4 Other MISRA Documents

- [16] MISRA *Development guidelines for vehicle based software* ISBN 0-9524156-0-7, Motor Industry Research Association, Nuneaton, November 1994
- [17] MISRA Report 5 *Software Metrics* Motor Industry Research Association, Nuneaton, February 1995
- [18] MISRA Report 6 *Verification and Validation* Motor Industry Research Association, Nuneaton, February 1995

9.2 The C Standard

- [19] ISO/IEC 9899:1990, *Programming languages — C* International Organization for Standardization, 1990
- [20] ISO/IEC 9899:1990/COR 1:1995, *Technical Corrigendum 1* International Organization for Standardization, 1995
- [21] ISO/IEC 9899:1990/AMD 1:1995, *Amendment 1* International Organization for Standardization, 1995
- [22] ISO/IEC 9899:1990/COR 2:1996, *Technical Corrigendum 2* International Organization for Standardization, 1996
- [23] ISO/IEC 9899:1999, *Programming languages — C* International Organization for Standardization, 1999
- [24] ISO/IEC 9899:1999/COR 1:2001, *Technical Corrigendum 1* International Organization for Standardization, 2001
- [25] ISO/IEC 9899:1999/COR 2:2004, *Technical Corrigendum 2* International Organization for Standardization, 2004
- [26] ISO/IEC 9899:1999/COR 3:2007, *Technical Corrigendum 3* International Organization for Standardization, 2007
- [27] ISO/IEC 9899:2011, *Programming languages — C* International Organization for Standardization, 2011
- [28] ISO/IEC 9899:2011/COR 1:2012, *Technical Corrigendum 1* International Organization for Standardization, 2012
- [29] ISO/IEC 9899:2018, *Programming languages — C* International Organization for Standardization, 2018

9.3 Other International Standards

- [30] ISO 9001:2015, *Quality management systems — Requirements*
International Organization for Standardization, 2015
- [31] ISO 90003:2018, *Software engineering — Guidelines for the application of ISO 9001:2015 to computer software*
International Organization for Standardization, 2018
- [32] ISO/IEC 10646:2020, *Information technology — Universal Multiple-Octet Coded Character Set (UCS)*
International Organization for Standardization, 2020
- [33] ISO/IEC 17961:2013, *Information technology — Programming languages, their environments and system software interfaces — C secure coding rules*
International Organization for Standardization, 2013
- [34] ISO 26262:2018, *Road vehicles — Functional safety*
International Organization for Standardization, 2018
- [35] IEC 60559:1989, *Binary floating-point arithmetic for microprocessor systems*,
International Organization for Standardization, 1989
Note: Subsequent to the publication of ISO/IEC 9899:2018 [29], IEC 60559:2020 has been published.
- [36] IEC 61508:2010, *Functional safety of electrical/electronic/programmable electronic safety-related systems*
International Electromechanical Commission, in 7 parts published in 2010
- [37] IEC 62304:2006, *Medical device software — Software life cycle processes*
International Electromechanical Commission, 2006
- [38] ARINC 653, *Avionics Application Software Standard Interface*,
Aeronautical Radio Inc., <https://aviation-ia.sae-itc.com/standards/>
- [39] DO-178C/ED-12C, *Software Considerations in Airborne Systems and Equipment Certification*
Radio Technical Commission for Aeronautics, 2011
- [40] EN 50128:2011, *Railway applications — Communications, signalling and processing systems — Software for railway control and protection*
CENELEC, 2011
- [41] RFC 3986, *Uniform Resource Identifier (URI): Generic Syntax*,
The Internet Society, 2005
Available from <https://www.ietf.org/rfc/rfc3986.txt>

9.4 Other References

- [42] AUTomotive Open System ARchitecture (AUTOSAR), <https://www.autosar.org>
- [43] CRR80, *The Use of Commercial Off-the-Shelf (COTS) Software in Safety Related Applications*
ISBN 0-7176-0984-7, HSE Books
- [44] Fenton N.E. and Pfleeger S.L., *Software Metrics: A Rigorous and Practical Approach*, 2nd Edition
ISBN 0-534-95429-1, PWS, 1998

- [45] Goldberg D., *What Every Computer Scientist Should Know about Floating-Point Arithmetic* Computing Surveys, March 1991
- [46] Hatton L., *Safer C: Developing Software for High-integrity and Safety-critical Systems* ISBN 0-07-707640-0, McGraw-Hill, 1994
- [47] Kernighan B.W., Ritchie D.M., *The C programming language*, 2nd edition ISBN 0-13-110362-8, Prentice Hall, 1988
(note: The 1st edition is not a suitable reference document as it does not describe ANSI/ISO C)
- [48] Koenig A., *C Traps and Pitfalls* ISBN 0-201-17928-8, Addison-Wesley, 1988
- [49] *OSEK/VDX Operating System*, Version 2.2.3., 2005
- [50] Software Engineering Center, Information-technology Promotion Agency, Japan (IPA/SEC), *Embedded System development Coding Reference (ESCR) [C language edition]* Version 3.0 SEC Books, Spring 2018
- [51] Straker D., *C Style: Standards and Guidelines* ISBN 0-13-116898-3, Prentice Hall, 1991
- [52] The TickIT Guide, *Using ISO 9001:2000 for Software Quality Management System Construction, Certification and Continual Improvement*, Issue 5 British Standards Institution, 2001

Appendix A Summary of guidelines

The implementation

Dir 1.1 Required Any implementation-defined behaviour on which the output of the program depends shall be documented and understood

Compilation and build

Dir 2.1 Required All source files shall compile without any compilation errors

Requirements traceability

Dir 3.1 Required All code shall be traceable to documented requirements

Code design

Dir 4.1 Required Run-time failures shall be minimized

Dir 4.2 Advisory All usage of assembly language should be documented

Dir 4.3 Required Assembly language shall be encapsulated and isolated

Dir 4.4 Advisory Sections of code should not be “commented out”

Dir 4.5 Advisory Identifiers in the same name space with overlapping visibility should be typographically unambiguous

Dir 4.6 Advisory typedefs that indicate size and signedness should be used in place of the basic numerical types

Dir 4.7 Required If a function returns error information, then that error information shall be tested

Dir 4.8 Advisory If a pointer to a structure or union is never dereferenced within a translation unit, then the implementation of the object should be hidden

Dir 4.9 Advisory A function should be used in preference to a function-like macro where they are interchangeable

Dir 4.10 Required Precautions shall be taken in order to prevent the contents of a header file being included more than once

Dir 4.11 Required The validity of values passed to library functions shall be checked

Dir 4.12 Required Dynamic memory allocation shall not be used

Dir 4.13 Advisory Functions which are designed to provide operations on a resource should be called in an appropriate sequence

Dir 4.14 Required The validity of values received from external sources shall be checked

Dir 4.15	Required	Evaluation of floating-point expressions shall not lead to the undetected generation of infinities and NaNs
----------	----------	---

Concurrency considerations

Dir 5.1	Required	There shall be no data races between threads
Dir 5.2	Required	There shall be no deadlocks between threads
Dir 5.3	Required	There shall be no dynamic thread creation

A standard C environment

Rule 1.1	Required	The program shall contain no violations of the standard C syntax and constraints, and shall not exceed the implementation's translation limits
Rule 1.2	Advisory	Language extensions should not be used
Rule 1.3	Required	There shall be no occurrence of undefined or critical unspecified behaviour
Rule 1.4	Required	Emergent language features shall not be used
Rule 1.5	Required	Obsolescent language features shall not be used

Unused code

Rule 2.1	Required	A project shall not contain unreachable code
Rule 2.2	Required	A project shall not contain dead code
Rule 2.3	Advisory	A project should not contain unused type declarations
Rule 2.4	Advisory	A project should not contain unused tag declarations
Rule 2.5	Advisory	A project should not contain unused macro definitions
Rule 2.6	Advisory	A function should not contain unused label declarations
Rule 2.7	Advisory	A function should not contain unused parameters
Rule 2.8	Advisory	A project should not contain unused object definitions

Comments

Rule 3.1	Required	The character sequences <code>/*</code> and <code>//</code> shall not be used within a comment
Rule 3.2	Required	Line-splicing shall not be used in <code>//</code> comments

Character sets and lexical conventions

Rule 4.1	Required	Octal and hexadecimal escape sequences shall be terminated
----------	----------	--

Rule 4.2 Advisory Trigraphs should not be used

Identifiers

Rule 5.1 Required External identifiers shall be distinct

Rule 5.2 Required Identifiers declared in the same scope and name space shall be distinct

Rule 5.3 Required An identifier declared in an inner scope shall not hide an identifier declared in an outer scope

Rule 5.4 Required Macro identifiers shall be distinct

Rule 5.5 Required Identifiers shall be distinct from macro names

Rule 5.6 Required A typedef name shall be a unique identifier

Rule 5.7 Required A tag name shall be a unique identifier

Rule 5.8 Required Identifiers that define objects or functions with external linkage shall be unique

Rule 5.9 Advisory Identifiers that define objects or functions with internal linkage should be unique

Types

Rule 6.1 Required Bit-fields shall only be declared with an appropriate type

Rule 6.2 Required Single-bit named bit-fields shall not be of a signed type

Rule 6.3 Required A bit field shall not be declared as a member of a union

Literals and constants

Rule 7.1 Required Octal constants shall not be used

Rule 7.2 Required A “u” or “U” suffix shall be applied to all integer constants that are represented in an unsigned type

Rule 7.3 Required The lowercase character “l” shall not be used in a literal suffix

Rule 7.4 Required A string literal shall not be assigned to an object unless the object’s type is “pointer to const-qualified char”

Rule 7.5 Mandatory The argument of an integer constant macro shall have an appropriate form

Rule 7.6 Required The small integer variants of the minimum-width integer constant macros shall not be used

Declarations and definitions

Rule 8.1 Required Types shall be explicitly specified

Rule 8.2	Required	Function types shall be in prototype form with named parameters
Rule 8.3	Required	All declarations of an object or function shall use the same names and type qualifiers
Rule 8.4	Required	A compatible declaration shall be visible when an object or function with external linkage is defined
Rule 8.5	Required	An external object or function shall be declared once in one and only one file
Rule 8.6	Required	An identifier with external linkage shall have exactly one external definition
Rule 8.7	Advisory	Functions and objects should not be defined with external linkage if they are referenced in only one translation unit
Rule 8.8	Required	The static storage class specifier shall be used in all declarations of objects and functions that have internal linkage
Rule 8.9	Advisory	An object should be declared at block scope if its identifier only appears in a single function
Rule 8.10	Required	An inline function shall be declared with the static storage class
Rule 8.11	Advisory	When an array with external linkage is declared, its size should be explicitly specified
Rule 8.12	Required	Within an enumerator list, the value of an implicitly-specified enumeration constant shall be unique
Rule 8.13	Advisory	A pointer should point to a const-qualified type whenever possible
Rule 8.14	Required	The restrict type qualifier shall not be used
Rule 8.15	Required	All declarations of an object with an explicit alignment specification shall specify the same alignment
Rule 8.16	Advisory	The alignment specification of zero should not appear in an object declaration
Rule 8.17	Advisory	At most one explicit alignment specifier should appear in an object declaration

Initialization

Rule 9.1	Mandatory	The value of an object with automatic storage duration shall not be read before it has been set
Rule 9.2	Required	The initializer for an aggregate or union shall be enclosed in braces
Rule 9.3	Required	Arrays shall not be partially initialized
Rule 9.4	Required	An element of an object shall not be initialized more than once

Rule 9.5	Required	Where designated initializers are used to initialize an array object the size of the array shall be specified explicitly
Rule 9.6	Required	An initializer using chained designators shall not contain initializers without designators
Rule 9.7	Mandatory	Atomic objects shall be appropriately initialized before being accessed

The essential type model

Rule 10.1	Required	Operands shall not be of an inappropriate essential type
Rule 10.2	Required	Expressions of essentially character type shall not be used inappropriately in addition and subtraction operations
Rule 10.3	Required	The value of an expression shall not be assigned to an object with a narrower essential type or of a different essential type category
Rule 10.4	Required	Both operands of an operator in which the usual arithmetic conversions are performed shall have the same essential type category
Rule 10.5	Advisory	The value of an expression should not be cast to an inappropriate essential type
Rule 10.6	Required	The value of a composite expression shall not be assigned to an object with wider essential type
Rule 10.7	Required	If a composite expression is used as one operand of an operator in which the usual arithmetic conversions are performed then the other operand shall not have wider essential type
Rule 10.8	Required	The value of a composite expression shall not be cast to a different essential type category or a wider essential type

Pointer type conversions

Rule 11.1	Required	Conversions shall not be performed between a pointer to a function and any other type
Rule 11.2	Required	Conversions shall not be performed between a pointer to an incomplete type and any other type
Rule 11.3	Required	A conversion shall not be performed between a pointer to object type and a pointer to a different object type
Rule 11.4	Advisory	A conversion should not be performed between a pointer to object and an integer type
Rule 11.5	Advisory	A conversion should not be performed from pointer to void into pointer to object
Rule 11.6	Required	A cast shall not be performed between pointer to void and an arithmetic type

Rule 11.7	Required	A cast shall not be performed between pointer to object and a non-integer arithmetic type
Rule 11.8	Required	A conversion shall not remove any const, volatile or <code>_Atomic</code> qualification from the type pointed to by a pointer
Rule 11.9	Required	The macro <code>NULL</code> shall be the only permitted form of integer null pointer constant
Rule 11.10	Required	The <code>_Atomic</code> qualifier shall not be applied to the incomplete type <code>void</code>

Expressions

Rule 12.1	Advisory	The precedence of operators within expressions should be made explicit
Rule 12.2	Required	The right hand operand of a shift operator shall lie in the range zero to one less than the width in bits of the essential type of the left hand operand
Rule 12.3	Advisory	The comma operator should not be used
Rule 12.4	Advisory	Evaluation of constant expressions should not lead to unsigned integer wrap-around
Rule 12.5	Mandatory	The <code>sizeof</code> operator shall not have an operand which is a function parameter declared as "array of type"
Rule 12.6	Required	Structure and union members of atomic objects shall not be directly accessed

Side effects

Rule 13.1	Required	Initializer lists shall not contain persistent side effects
Rule 13.2	Required	The value of an expression and its persistent side effects shall be the same under all permitted evaluation orders and shall be independent from thread interleaving
Rule 13.3	Advisory	A full expression containing an increment (<code>++</code>) or decrement (<code>--</code>) operator should have no other potential side effects other than that caused by the increment or decrement operator
Rule 13.4	Advisory	The result of an assignment operator should not be used
Rule 13.5	Required	The right hand operand of a logical <code>&&</code> or <code> </code> operator shall not contain persistent side effects
Rule 13.6	Required	The operand of the <code>sizeof</code> operator shall not contain any expression which has potential side effects

Control statement expressions

Rule 14.1	Required	A loop counter shall not have essentially floating type
Rule 14.2	Required	A for loop shall be well-formed

Rule 14.3	Required	Controlling expressions shall not be invariant
Rule 14.4	Required	The controlling expression of an if statement and the controlling expression of an iteration-statement shall have essentially Boolean type

Control flow

Rule 15.1	Advisory	The goto statement should not be used
Rule 15.2	Required	The goto statement shall jump to a label declared later in the same function
Rule 15.3	Required	Any label referenced by a goto statement shall be declared in the same block, or in any block enclosing the goto statement
Rule 15.4	Advisory	There should be no more than one break or goto statement used to terminate any iteration statement
Rule 15.5	Advisory	A function should have a single point of exit at the end
Rule 15.6	Required	The body of an iteration-statement or a selection-statement shall be a compound-statement
Rule 15.7	Required	All if ... else if constructs shall be terminated with an else statement

Switch statements

Rule 16.1	Required	All switch statements shall be well-formed
Rule 16.2	Required	A switch label shall only be used when the most closely-enclosing compound statement is the body of a switch statement
Rule 16.3	Required	An unconditional break statement shall terminate every switch-clause
Rule 16.4	Required	Every switch statement shall have a default label
Rule 16.5	Required	A default label shall appear as either the first or the last switch label of a switch statement
Rule 16.6	Required	Every switch statement shall have at least two switch-clauses
Rule 16.7	Required	A switch-expression shall not have essentially Boolean type

Functions

Rule 17.1	Required	The standard header file <stdarg.h> shall not be used
Rule 17.2	Required	Functions shall not call themselves, either directly or indirectly
Rule 17.3	Mandatory	A function shall not be declared implicitly
Rule 17.4	Mandatory	All exit paths from a function with non-void return type shall have an explicit return statement with an expression

Rule 17.5	Required	The function argument corresponding to a parameter declared to have an array type shall have an appropriate number of elements
Rule 17.6	Mandatory	The declaration of an array parameter shall not contain the static keyword between the []
Rule 17.7	Required	The value returned by a function having non-void return type shall be used
Rule 17.8	Advisory	A function parameter should not be modified
Rule 17.9	Mandatory	A function declared with a <code>_Noreturn</code> function specifier shall not return to its caller
Rule 17.10	Required	A function declared with a <code>_Noreturn</code> function specifier shall have void return type
Rule 17.11	Advisory	A function that never returns should be declared with a <code>_Noreturn</code> function specifier
Rule 17.12	Advisory	A function identifier should only be used with either a preceding <code>&</code> , or with a parenthesized parameter list
Rule 17.13	Required	A function type shall not be type qualified

Pointers and arrays

Rule 18.1	Required	A pointer resulting from arithmetic on a pointer operand shall address an element of the same array as that pointer operand
Rule 18.2	Required	Subtraction between pointers shall only be applied to pointers that address elements of the same array
Rule 18.3	Required	The relational operators <code>></code> , <code>>=</code> , <code><</code> and <code><=</code> shall not be applied to expressions of pointer type except where they point into the same object
Rule 18.4	Advisory	The <code>+</code> , <code>-</code> , <code>+=</code> and <code>-=</code> operators should not be applied to an expression of pointer type
Rule 18.5	Advisory	Declarations should contain no more than two levels of pointer nesting
Rule 18.6	Required	The address of an object with automatic or thread-local storage shall not be copied to another object that persists after the first object has ceased to exist
Rule 18.7	Required	Flexible array members shall not be declared
Rule 18.8	Required	Variable-length arrays shall not be used
Rule 18.9	Required	An object with temporary lifetime shall not undergo array-to-pointer conversion
Rule 18.10	Mandatory	Pointers to variably-modified array types shall not be used

Overlapping storage

Rule 19.1	Mandatory	An object shall not be assigned or copied to an overlapping object
Rule 19.2	Advisory	The union keyword should not be used

Preprocessing directives

Rule 20.1	Advisory	<code>#include</code> directives should only be preceded by preprocessor directives or comments
Rule 20.2	Required	The <code>'</code> , <code>"</code> or <code>\</code> characters and the <code>/*</code> or <code>//</code> character sequences shall not occur in a header file name
Rule 20.3	Required	The <code>#include</code> directive shall be followed by either a <code><filename></code> or <code>"filename"</code> sequence
Rule 20.4	Required	A macro shall not be defined with the same name as a keyword
Rule 20.5	Advisory	<code>#undef</code> should not be used
Rule 20.6	Required	Tokens that look like a preprocessing directive shall not occur within a macro argument
Rule 20.7	Required	Expressions resulting from the expansion of macro parameters shall be enclosed in parentheses
Rule 20.8	Required	The controlling expression of a <code>#if</code> or <code>#elif</code> preprocessing directive shall evaluate to 0 or 1
Rule 20.9	Required	All identifiers used in the controlling expression of <code>#if</code> or <code>#elif</code> preprocessing directives shall be <code>#define</code> 'd before evaluation
Rule 20.10	Advisory	The <code>#</code> and <code>##</code> preprocessor operators should not be used
Rule 20.11	Required	A macro parameter immediately following a <code>#</code> operator shall not immediately be followed by a <code>##</code> operator
Rule 20.12	Required	A macro parameter used as an operand to the <code>#</code> or <code>##</code> operators, which is itself subject to further macro replacement, shall only be used as an operand to these operators
Rule 20.13	Required	A line whose first token is <code>#</code> shall be a valid preprocessing directive
Rule 20.14	Required	All <code>#else</code> , <code>#elif</code> and <code>#endif</code> preprocessor directives shall reside in the same file as the <code>#if</code> , <code>#ifdef</code> or <code>#ifndef</code> directive to which they are related

Standard libraries

Rule 21.1	Required	<code>#define</code> and <code>#undef</code> shall not be used on a reserved identifier or reserved macro name
Rule 21.2	Required	A reserved identifier or reserved macro name shall not be declared

Rule 21.3	Required	The memory allocation and deallocation functions of <stdlib.h> shall not be used
Rule 21.4	Required	The standard header file <setjmp.h> shall not be used
Rule 21.5	Required	The standard header file <signal.h> shall not be used
Rule 21.6	Required	The Standard Library input/output functions shall not be used
Rule 21.7	Required	The Standard Library functions atof, atoi, atol and atoll of <stdlib.h> shall not be used
Rule 21.8	Required	The Standard Library termination functions of <stdlib.h> shall not be used
Rule 21.9	Required	The Standard Library functions bsearch and qsort of <stdlib.h> shall not be used
Rule 21.10	Required	The Standard Library time and date functions shall not be used
Rule 21.11	Advisory	The standard header file <tgmath.h> should not be used
Rule 21.12	Required	The standard header file <fenv.h> shall not be used
Rule 21.13	Mandatory	Any value passed to a function in <ctype.h> shall be representable as an unsigned char or be the value EOF
Rule 21.14	Required	The Standard Library function memcmp shall not be used to compare null terminated strings
Rule 21.15	Required	The pointer arguments to the Standard Library functions memcpy, memmove and memcmp shall be pointers to qualified or unqualified versions of compatible types
Rule 21.16	Required	The pointer arguments to the Standard Library function memcmp shall point to either a pointer type, an essentially signed type, an essentially unsigned type, an essentially Boolean type or an essentially enum type
Rule 21.17	Mandatory	Use of the string handling functions from <string.h> shall not result in accesses beyond the bounds of the objects referenced by their pointer parameters
Rule 21.18	Mandatory	The size_t argument passed to any function in <string.h> shall have an appropriate value
Rule 21.19	Mandatory	The pointers returned by the Standard Library functions localeconv, getenv, setlocale or, strerror shall only be used as if they have pointer to const-qualified type
Rule 21.20	Mandatory	The pointer returned by the Standard Library functions asctime, ctime, gmtime, localtime, localeconv, getenv, setlocale or strerror shall not be used following a subsequent call to the same function
Rule 21.21	Required	The Standard Library function system of <stdlib.h> shall not be used
Rule 21.22	Mandatory	All operand arguments to any type-generic macros declared in <tgmath.h> shall have an appropriate essential type

Rule 21.23	Required	All operand arguments to any multi-argument type-generic macros declared in <code><tgmath.h></code> shall have the same standard type
Rule 21.24	Required	The random number generator functions of <code><stdlib.h></code> shall not be used
Rule 21.25	Required	All memory synchronization operations shall be executed in sequentially consistent order
Rule 21.26	Required	The Standard Library function <code>mtx_timedlock()</code> shall only be invoked on mutex objects of appropriate mutex type

Resources

Rule 22.1	Required	All resources obtained dynamically by means of Standard Library functions shall be explicitly released
Rule 22.2	Mandatory	A block of memory shall only be freed if it was allocated by means of a Standard Library function
Rule 22.3	Required	The same file shall not be open for read and write access at the same time on different streams
Rule 22.4	Mandatory	There shall be no attempt to write to a stream which has been opened as read-only
Rule 22.5	Mandatory	A pointer to a FILE object shall not be dereferenced
Rule 22.6	Mandatory	The value of a pointer to a FILE shall not be used after the associated stream has been closed
Rule 22.7	Required	The macro <code>EOF</code> shall only be compared with the unmodified return value from any Standard Library function capable of returning <code>EOF</code>
Rule 22.8	Required	The value of <code>errno</code> shall be set to zero prior to a call to an <code>errno</code> -setting-function
Rule 22.9	Required	The value of <code>errno</code> shall be tested against zero after calling an <code>errno</code> -setting-function
Rule 22.10	Required	The value of <code>errno</code> shall only be tested when the last function to be called was an <code>errno</code> -setting-function
Rule 22.11	Required	A thread that was previously either joined or detached shall not be subsequently joined nor detached
Rule 22.12	Mandatory	Thread objects, thread synchronization objects, and thread-specific storage pointers shall only be accessed by the appropriate Standard Library functions
Rule 22.13	Required	Thread objects, thread synchronization objects and thread-specific storage pointers shall have appropriate storage duration
Rule 22.14	Mandatory	Thread synchronization objects shall be initialized before being accessed

Rule 22.15	Required	Thread synchronization objects and thread-specific storage pointers shall not be destroyed until after all threads accessing them have terminated
Rule 22.16	Required	All mutex objects locked by a thread shall be explicitly unlocked by the same thread
Rule 22.17	Required	No thread shall unlock a mutex or call <code>cond_wait()</code> or <code>cond_timedwait()</code> for a mutex it has not locked before
Rule 22.18	Required	Non-recursive mutexes shall not be recursively locked
Rule 22.19	Required	A condition variable shall be associated with at most one mutex object
Rule 22.20	Mandatory	Thread-specific storage pointers shall be created before being accessed

Generic Selections

Rule 23.1	Advisory	A generic selection should only be expanded from a macro
Rule 23.2	Required	A generic selection that is not expanded from a macro shall not contain potential side effects in the controlling expression
Rule 23.3	Advisory	A generic selection should contain at least one non-default association
Rule 23.4	Required	A generic association shall list an appropriate type
Rule 23.5	Advisory	A generic selection should not depend on implicit pointer type conversion
Rule 23.6	Required	The controlling expression of a generic selection shall have an essential type that matches its standard type
Rule 23.7	Advisory	A generic selection that is expanded from a macro should evaluate its argument only once
Rule 23.8	Required	A default association shall appear as either the first or the last association of a generic selection

Appendix B Guideline attributes

Rule	Category	Applies to	Analysis
Dir 1.1	Required	C90, C99, C11	
Dir 2.1	Required	C90, C99, C11	
Dir 3.1	Required	C90, C99, C11	
Dir 4.1	Required	C90, C99, C11	
Dir 4.2	Advisory	C90, C99, C11	
Dir 4.3	Required	C90, C99, C11	
Dir 4.4	Advisory	C90, C99, C11	
Dir 4.5	Advisory	C90, C99, C11	
Dir 4.6	Advisory	C90, C99, C11	
Dir 4.7	Required	C90, C99, C11	
Dir 4.8	Advisory	C90, C99, C11	
Dir 4.9	Advisory	C90, C99, C11	
Dir 4.10	Required	C90, C99, C11	
Dir 4.11	Required	C90, C99, C11	
Dir 4.12	Required	C90, C99, C11	
Dir 4.13	Advisory	C90, C99, C11	
Dir 4.14	Required	C90, C99, C11	
Dir 4.15	Required	C90, C99, C11	
Dir 5.1	Required	C11	
Dir 5.2	Required	C11	
Dir 5.3	Required	C11	
Rule 1.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 1.2	Advisory	C90, C99, C11	Undecidable, Single Translation Unit
Rule 1.3	Required	C90, C99, C11	Undecidable, System
Rule 1.4	Required	C11	Decidable, Single Translation Unit
Rule 1.5	Required	C99, C11	Undecidable, System
Rule 2.1	Required	C90, C99, C11	Undecidable, System
Rule 2.2	Required	C90, C99, C11	Undecidable, System
Rule 2.3	Advisory	C90, C99, C11	Decidable, System
Rule 2.4	Advisory	C90, C99, C11	Decidable, System
Rule 2.5	Advisory	C90, C99, C11	Decidable, System
Rule 2.6	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 2.7	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 2.8	Advisory	C90, C99, C11	Decidable, System
Rule 3.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 3.2	Required	C99, C11	Decidable, Single Translation Unit
Rule 4.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 4.2	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 5.1	Required	C90, C99, C11	Decidable, System
Rule 5.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 5.3	Required	C90, C99, C11	Decidable, Single Translation Unit

Rule	Category	Applies to	Analysis
Rule 5.4	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 5.5	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 5.6	Required	C90, C99, C11	Decidable, System
Rule 5.7	Required	C90, C99, C11	Decidable, System
Rule 5.8	Required	C90, C99, C11	Decidable, System
Rule 5.9	Advisory	C90, C99, C11	Decidable, System
Rule 6.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 6.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 6.3	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 7.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 7.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 7.3	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 7.4	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 7.5	Mandatory	C99, C11	Decidable, Single Translation Unit
Rule 7.6	Required	C99, C11	Decidable, Single Translation Unit
Rule 8.1	Required	C90	Decidable, Single Translation Unit
Rule 8.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 8.3	Required	C90, C99, C11	Decidable, System
Rule 8.4	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 8.5	Required	C90, C99, C11	Decidable, System
Rule 8.6	Required	C90, C99, C11	Decidable, System
Rule 8.7	Advisory	C90, C99, C11	Decidable, System
Rule 8.8	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 8.9	Advisory	C90, C99, C11	Decidable, System
Rule 8.10	Required	C99, C11	Decidable, Single Translation Unit
Rule 8.11	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 8.12	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 8.13	Advisory	C90, C99, C11	Undecidable, System
Rule 8.14	Required	C99, C11	Decidable, Single Translation Unit
Rule 8.15	Required	C11	Decidable, System
Rule 8.16	Advisory	C11	Decidable, Single Translation Unit
Rule 8.17	Advisory	C11	Decidable, Single Translation Unit
Rule 9.1	Mandatory	C90, C99, C11	Undecidable, System
Rule 9.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 9.3	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 9.4	Required	C99, C11	Decidable, Single Translation Unit
Rule 9.5	Required	C99, C11	Decidable, Single Translation Unit
Rule 9.6	Required	C99, C11	Decidable, Single Translation Unit
Rule 9.7	Mandatory	C11	Undecidable, System
Rule 10.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 10.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 10.3	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 10.4	Required	C90, C99, C11	Decidable, Single Translation Unit

Rule	Category	Applies to	Analysis
Rule 10.5	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 10.6	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 10.7	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 10.8	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.3	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.4	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.5	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.6	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.7	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.8	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.9	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 11.10	Required	C11	Decidable, Single Translation Unit
Rule 12.1	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 12.2	Required	C90, C99, C11	Undecidable, System
Rule 12.3	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 12.4	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 12.5	Mandatory	C90, C99, C11	Decidable, Single Translation Unit
Rule 12.6	Required	C11	Decidable, Single Translation Unit
Rule 13.1	Required	C99, C11	Undecidable, System
Rule 13.2	Required	C90, C99, C11	Undecidable, System
Rule 13.3	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 13.4	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 13.5	Required	C90, C99, C11	Undecidable, System
Rule 13.6	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 14.1	Required	C90, C99, C11	Undecidable, System
Rule 14.2	Required	C90, C99, C11	Undecidable, System
Rule 14.3	Required	C90, C99, C11	Undecidable, System
Rule 14.4	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 15.1	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 15.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 15.3	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 15.4	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 15.5	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 15.6	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 15.7	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 16.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 16.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 16.3	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 16.4	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 16.5	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 16.6	Required	C90, C99, C11	Decidable, Single Translation Unit

Rule	Category	Applies to	Analysis
Rule 16.7	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 17.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 17.2	Required	C90, C99, C11	Undecidable, System
Rule 17.3	Mandatory	C90	Decidable, Single Translation Unit
Rule 17.4	Mandatory	C90, C99, C11	Decidable, Single Translation Unit
Rule 17.5	Required	C90, C99, C11	Undecidable, System
Rule 17.6	Mandatory	C99, C11	Decidable, Single Translation Unit
Rule 17.7	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 17.8	Advisory	C90, C99, C11	Undecidable, System
Rule 17.9	Mandatory	C11	Undecidable, System
Rule 17.10	Required	C11	Decidable, Single Translation Unit
Rule 17.11	Advisory	C11	Undecidable, System
Rule 17.12	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 17.13	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 18.1	Required	C90, C99, C11	Undecidable, System
Rule 18.2	Required	C90, C99, C11	Undecidable, System
Rule 18.3	Required	C90, C99, C11	Undecidable, System
Rule 18.4	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 18.5	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 18.6	Required	C90, C99, C11	Undecidable, System
Rule 18.7	Required	C99, C11	Decidable, Single Translation Unit
Rule 18.8	Required	C99, C11	Decidable, Single Translation Unit
Rule 18.9	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 18.10	Mandatory	C99, C11	Decidable, Single Translation Unit
Rule 19.1	Mandatory	C90, C99, C11	Undecidable, System
Rule 19.2	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.1	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.3	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.4	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.5	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.6	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.7	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.8	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.9	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.10	Advisory	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.11	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.12	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.13	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 20.14	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.1	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.2	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.3	Required	C90, C99, C11	Decidable, Single Translation Unit

Rule	Category	Applies to	Analysis
Rule 21.4	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.5	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.6	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.7	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.8	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.9	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.10	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.11	Advisory	C99, C11	Decidable, Single Translation Unit
Rule 21.12	Required	C99, C11	Decidable, Single Translation Unit
Rule 21.13	Mandatory	C90, C99, C11	Undecidable, System
Rule 21.14	Required	C90, C99, C11	Undecidable, System
Rule 21.15	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.16	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.17	Mandatory	C90, C99, C11	Undecidable, System
Rule 21.18	Mandatory	C90, C99, C11	Undecidable, System
Rule 21.19	Mandatory	C90, C99, C11	Undecidable, System
Rule 21.20	Mandatory	C90, C99, C11	Undecidable, System
Rule 21.21	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.22	Mandatory	C99, C11	Decidable, Single Translation Unit
Rule 21.23	Required	C99, C11	Decidable, Single Translation Unit
Rule 21.24	Required	C90, C99, C11	Decidable, Single Translation Unit
Rule 21.25	Required	C11	Decidable, Single Translation Unit
Rule 21.26	Required	C11	Undecidable, System
Rule 22.1	Required	C90, C99, C11	Undecidable, System
Rule 22.2	Mandatory	C90, C99, C11	Undecidable, System
Rule 22.3	Required	C90, C99, C11	Undecidable, System
Rule 22.4	Mandatory	C90, C99, C11	Undecidable, System
Rule 22.5	Mandatory	C90, C99, C11	Undecidable, System
Rule 22.6	Mandatory	C90, C99, C11	Undecidable, System
Rule 22.7	Required	C90, C99, C11	Undecidable, System
Rule 22.8	Required	C90, C99, C11	Undecidable, System
Rule 22.9	Required	C90, C99, C11	Undecidable, System
Rule 22.10	Required	C90, C99, C11	Undecidable, System
Rule 22.11	Required	C11	Undecidable, System
Rule 22.12	Mandatory	C11	Undecidable, System
Rule 22.13	Required	C11	Decidable, Single Translation Unit
Rule 22.14	Mandatory	C11	Undecidable, System
Rule 22.15	Required	C11	Undecidable, System
Rule 22.16	Required	C11	Undecidable, System
Rule 22.17	Required	C11	Undecidable, System
Rule 22.18	Required	C11	Undecidable, System
Rule 22.19	Required	C11	Undecidable, System
Rule 22.20	Mandatory	C11	Undecidable, System

Rule	Category	Applies to	Analysis
Rule 23.1	Advisory	C11	Decidable, Single Translation Unit
Rule 23.2	Required	C11	Decidable, Single Translation Unit
Rule 23.3	Advisory	C11	Decidable, Single Translation Unit
Rule 23.4	Required	C11	Decidable, Single Translation Unit
Rule 23.5	Advisory	C11	Decidable, Single Translation Unit
Rule 23.6	Required	C11	Decidable, Single Translation Unit
Rule 23.7	Advisory	C11	Decidable, Single Translation Unit
Rule 23.8	Required	C11	Decidable, Single Translation Unit

Appendix C Type safety issues with C

The C Standard may be considered to exhibit poor type safety as it permits a wide range of implicit type conversions to take place. These type conversions can compromise safety as their implementation-defined aspects can cause developer confusion.

Note: Conversions may be implicit or explicit (i.e. by means of a cast) — where the term *conversion* is used without qualification it means either or both forms as the situation requires.

An explanation of these conversions follows, along with some examples to show how they can lead to developer confusion.

C.1 Type conversions

C.1.1 Implicit conversions

Implicit type conversions within ISO C permit the use of mixed types within expressions and may occur even when the types used are all the same. There are three types of implicit conversion:

1. An *integer promotion* conversion to *signed int* or *unsigned int* takes place whenever a bit-field or object with *(un)signed char*, *(un)signed short* or *enum* type is used arithmetically as described in C90 Section 6.2.1.1, C99/C11 Section 6.3.1.1. *Integer promotion* preserves the original value, but the signedness may change (e.g. *unsigned char* will be converted to *signed int*);
2. The *usual arithmetic conversions* convert (balance) operands to a common type whenever operands with different types are used with certain operators, as described in C90 Section 6.2.1.5, C99/C11 Section 6.3.1.8;
3. Conversion on *assignment*.

An expression may contain none, one, two or all of these implicit conversions, as the examples below show (*si* and *ui* are 32-bit signed and unsigned integers and *u8* is an 8-bit *unsigned char*):

Expression	Promote	Balance	Convert	Notes
<code>si = si + si;</code>				
<code>si = uc + uc;</code>	Yes			<i>uc</i> promoted to <i>signed int</i>
<code>ui = si + ui;</code>		Yes		<i>si</i> is balanced to <i>unsigned int</i>
<code>ui = ui + uc;</code>	Yes	Yes		<i>uc</i> promoted to <i>signed int</i> and balanced to <i>unsigned int</i>
<code>ui = si;</code>			Yes	<i>si</i> converted to <i>unsigned int</i>
<code>ui = uc + uc;</code>	Yes		Yes	<i>uc</i> promoted to <i>signed int</i> and result of expression converted to <i>unsigned int</i>
<code>si = si + ui;</code>		Yes	Yes	<i>si</i> is balanced to <i>unsigned int</i> and result of expression converted to <i>signed int</i>
<code>si = ui + uc;</code>	Yes	Yes	Yes	<i>uc</i> promoted to <i>signed int</i> , balanced to <i>unsigned int</i> and result of expression converted to <i>signed int</i>

C.1.2 Explicit conversions

Explicit type conversions (casts) may be introduced for functional reasons:

- To change the type in which a subsequent arithmetic operation is performed;
- To truncate a value deliberately;
- To make a type conversion explicit in the interests of clarity.

ISO C does not require all explicit casts to be checked, allowing conversions to take place between incompatible types.

C.1.3 Concerns with conversions

Implicit and explicit conversion can lead to several concerns, including:

- **Loss of value:** e.g. conversion to a type where the magnitude of the value cannot be represented;
- **Loss of sign:** e.g. conversion from a signed type to an unsigned type resulting in loss of sign;
- **Loss of imaginary part:** e.g. conversion of a complex floating type to a real floating type, where the imaginary part of the complex value is discarded;
- **Loss of precision:** e.g. conversion from a floating type to an integer type with consequent loss of precision;
- **Loss of layout:** e.g. conversion from a pointer to one type to a pointer to a different type leading to an incompatible storage layout.

Whilst many type conversions are safe, the only ones that can be guaranteed safe for all data values and all possible conforming implementations are:

- Conversion of an integer value to a wider integer type of the same signedness;
- Conversion of a floating type to a wider floating type;
- Conversion of a real floating type to a complex floating type, having wider or equal *corresponding real type*.

Depending on the implemented integer sizes, other type conversions may also be safe. However, any potentially dangerous type conversions should be identified by making them explicit.

C.2 Developer confusion

As the sizes in which types are implemented can vary between implementations, the contexts in which implicit conversions take place can also vary. This places a requirement on the developer to maintain a mental model of the type system in use for any particular project. It is relatively easy for a developer to use the wrong model from time-to-time, especially if they are working on multiple projects having different implementations.

C.2.1 Type widening in integer promotion

The type in which integer expressions are evaluated depends on the type of the operands after any integer promotion. Multiplying two 8-bit values will always give a result that is at least 16 bits wide. However, multiplying two 16-bit values will only give a 32-bit result when the implemented size of *int* is at least 32 bits. It is safer never to rely on the type-widening obtained by integer promotion as the associated implementation-defined behaviour may lead to developer confusion. Consider the following example:

```
uint16_t u16a = 40000;          /* unsigned short or unsigned int ? */
uint16_t u16b = 30000;          /* unsigned short or unsigned int ? */
uint32_t u32x;                  /* unsigned int or unsigned long ? */

u32x = u16a + u16b;             /* u32x = 70000 or 4464 ?      */
```

The expected result is possibly 70 000, but the value assigned to `u32x` will depend on the implemented size of *int*. If this is 32 bits, the addition will occur in 32-bit signed arithmetic and the “correct” value will be obtained. If it is only 16 bits, the addition will take place in 16-bit unsigned arithmetic, wraparound will occur and will yield the value 4 464 (70 000 % 65 536). Wraparound in unsigned arithmetic is well defined but, as its effects depend on the type in which the arithmetic is performed, its intentional use should be documented.

C.2.2 Evaluation type confusion

A similar problem arises where a developer is deceived into thinking that the type in which a calculation is conducted is influenced by the type to which the result is assigned or converted. For example, in the following code the two 16-bit objects are added together in 16-bit arithmetic (assuming *int* is 16-bit), and the result is converted to type `uint32_t` on assignment.

```
u32x = u16a + u16b;
```

It is not unusual for developers to be deceived into thinking that the addition is performed in 32-bit arithmetic, because of the type of `u32x`.

Confusion of this nature is not confined to integer arithmetic or to implicit conversions. The following examples demonstrate some statements in which the result is well defined but the calculation may not be performed in the type that the developer expects:

```
u32a = (uint32_t)(u16a * u16b); /* May not be a 32-bit operation */
f64a = u16a / u16b;           /* Not floating point division */
f32a = (float32_t)(u16a / u16b); /* Not floating point division */
f64a = f32a + f32b;           /* May be a low precision operation */
f64a = (float64_t)(f32a + f32b); /* May be a low precision operation */
```

C.2.3 Change of signedness in arithmetic operations

Integer promotion may lead to the result of an expression not meeting developer expectations. For example, the expression `10 - u16a` yields a result that is dependent on the size of *int*. If `u16a` has a value of 100 and *int* is 32 bits, then the result is signed with a value of -90. However, if *int* is 16 bits the result will be unsigned with a value of 65 446.

C.2.4 Change of signedness in bitwise operations

Integer promotion arising when bitwise operators are applied to small unsigned types can lead to confusion. For example a bitwise complement operation on an operand of type *unsigned char* will generally yield a result of type (*signed*) *int* with a negative value. The operand is promoted to type *int* before the operation and the extra high order bits are set by the complement process. The number of extra bits, if any, is dependent on the size of an *int* and it is particularly undesirable if the complement operation is followed by a right shift as this leads to implementation-defined behaviour.

```
u8a = 0xff;
if ( ~u8a == 0x00U ) /* This test will always fail */
```

Appendix D Essential types

The C language as defined by the ISO C standard contains a number of weaknesses and inconsistencies and, unfortunately, the *standard type* of an expression is not always fully descriptive of the essential nature of the data being represented. For example:

- **Integer promotion:** The set of integer types do not behave in a consistent way. As a result of integer promotion, an expression which is intrinsically *unsigned* (e.g. *unsigned char + unsigned char*) typically has a *standard type* of *signed int*.
- **Integer constants:** Integer constants only exist for *signed/unsigned int*, *long* and *long long* types. This complicates the issue of ensuring type consistency (for example when passing constants as function arguments).
- **Character constants:** The *standard type* of a character constant (e.g. `'x'`) is defined to be *int* (not *char*). This obscures the distinction between character data (i.e. “characters”) and numeric data (i.e. “number values”).
- **Logical expressions:** A Boolean type does not exist in the C90 language. Type *_Bool* was introduced into C99, but unfortunately, for reasons of backwards compatibility, the result of equality (`==`, `!=`), relational (`<`, `<=`, `>=`, `>`), and logical (`&&`, `||`, `!`) operators is still of type *int* rather than type *_Bool*.
- **Bit-fields:** ISO C does not define a bit-field type. This means that it is not possible, for example, to cast to a bit-field type. The *standard type* of a bit-field as defined by ISO C is one of *_Bool*, *signed int* or *unsigned int*. The signedness of a bit-field of type *int* is implementation-defined.
- **Enumerations:** An enumeration has a type which is implementation-defined. The implementation is free to use any integer type which is capable of representing the enumeration. If the enumeration includes negative values, the type will be a signed integer type. If the enumeration consists only of non-negative values, the type may be either a signed or unsigned integer type.
- **Enumeration constants:** Regardless of the type used to implement an enumeration, an enumeration constant is always of type *int*.
- **Complex floating-point expressions:** There is no promotion from real floating to complex floating for the operands of a complex expression.

D.1 The essential type category of expressions

MISRA C:2012 introduced the concept of *essential type* to help mitigate the issues identified above.

The *essential type category* of an expression is one of:

- *Essentially Boolean*;
- *Essentially character*;
- *Essentially enum*;
- *Essentially signed*;
- *Essentially unsigned*;

- *Essentially floating*, which may be either:
 - *Essentially real floating*, or
 - *Essentially complex floating*.

Note: For the purposes of this *essential type model*, *essentially real floating* and *essentially complex floating* are considered to be distinct *essential type categories*.

The following table shows how the standard integer types map on to *essential type categories*.

<i>Essential type category</i>				
Boolean	character	signed	unsigned	enum<i>
_Bool	char	signed char signed short signed int signed long signed long long	unsigned char unsigned short unsigned int unsigned long unsigned long long	named enum

<i>Essential type category</i>	
floating	
real floating	complex floating
float double long double	float _Complex double _Complex long double _Complex

Note: for C99 and later, any *extended integer type* is allocated a location appropriate to the rank and signedness inferred by the C Standard.

The concept of *rank* is particularly relevant when describing the set of signed and unsigned types; *signed* and *unsigned long long* share the “highest” rank and *signed* and *unsigned char* share the “lowest” rank. In C99 and later, “rank” is a term applied only to integer types.

The *essential type* of an expression only differs from the standard C type (*standard type*) in expressions where the *standard type* is either *signed int* or *unsigned int*.

The *essential type* of an expression with a C type defined in a C Standard Library header is that in which it is implemented. For example, `int_least8_t` may be implemented as *signed int* and will have an essential type of *essentially signed int* in that case.

In the following paragraphs the specific situations are defined where *essential type* and *standard type* are different.

D.2 The essential type of character data

An object with a *standard type* of *char* (i.e. single byte character) has *essentially char type*.

D.3 The signed and unsigned type of lowest rank (STLR and UTLR)

The *Essential Type Model* introduces the concept of the *Signed Type of Lowest Rank* (the *STLR*) and the *Unsigned Type of Lowest Rank* (the *UTLR*). These allow *integer constant expressions* and bit-fields to be used within expressions having an *essential type* with a rank lower than that of *int*. In the case of *integer constant expressions*, this is a convenience as it avoids the need to cast the expression, or its operands, to a type with lower rank. Similarly, for an implementation that does not permit bit-fields to have a type whose rank is lower than that of *int*, it avoids the need to cast bit-fields in some situations.

1. The *STLR* is the *signed type* having the lowest rank required to represent the value of a particular *integer constant expression* or both the maximum and minimum values of a bit-field;
2. The *UTLR* is the *unsigned type* having the lowest rank required to represent the value of a particular *integer constant expression* or the maximum value of a bit-field.

Note: The *STLR* and *UTLR* of an *integer constant expression* is only applied to those operators listed in Appendix D.7.

D.4 The essential type of bit-fields

The *essential type* of a bit-field is determined by the first of the following that applies:

1. For a bit-field which is implemented with an *essentially Boolean type* it is *essentially Boolean*;
2. For a bit-field which is implemented with a *signed type* it is the *STLR* which is able to represent the bit-field;
3. For a bit-field which is implemented with an *unsigned type* it is the *UTLR* which is able to represent the bit-field.

D.5 The essential type of enumerations

Two distinct types of enumeration need to be considered:

1. A *named enum type* is an enumeration which has a *tag* or which is used in the definition of any object, function or type;
2. An *anonymous enum type* is an enumeration which does not have a *tag* and which is not used in the definition of any object, function or type. This will typically be used to define a set of constants, which may or may not be related.

A *named enum type* is distinct from all other *named enum types*, even if declared in an inner scope with exactly the same tag and enumeration constants. Each instance of a *named enum type* is denoted as *enum<i>* in this document, with the *i* being different for each such type.

An alias for a *named enum type* created using *typedef* denotes that *named enum type* and is not a new type.

The *essential type* of an *anonymous enum type* is its *standard type*.

The following all have *named enum type* and have distinct *essential types*:

```
enum ETAG { A, B, C };
typedef enum { A, B, C } ETYPE;
typedef enum ETAG { A, B, C } ETYPE;
enum { A, B, C } x;
```

The following *typedefs* both refer to the same *enum<i>*:

```
typedef enum      { A, B, C } ETYPE;
typedef ETYPE    FTYPE;
```

D.6 The essential type of literal constants

The C Standard defines the following constants of *integer type*:

- Integer constants;
- Enumeration constants;
- Character constants.

Note: a constant of *integer type* is not necessarily an *integer constant*.

D.6.1 Integer constants

1. If the *standard type* of an integer constant is *signed int* then its *essential type* is the *STLR*;
2. If the *standard type* of an integer constant is *unsigned int* then its *essential type* is the *UTLR*.

D.6.2 Enumeration constants

The *standard type* of an enumeration constant in C is always *int*, regardless of the implemented type of the enumeration and regardless of the type of any initializer expression used to define its value. For example, if an enumeration constant is initialized with an unsigned value (e.g. 500u), the constant will still be considered to have a *standard type* of *signed int*.

The *essential type* of an enumeration constant is determined as follows:

1. If an enumeration defines a *named enum type* then the *essential type* of its enumeration constants is *enum<i>*;
 - 1.1 If a *named enum type* is used to define an *essentially Boolean type* then the *essential type* of its enumeration constants is *essentially Boolean*.
2. If an enumeration defines an *anonymous enum type* then the *essential type* of each enumeration constant is the *STLR* of its value.

In the following, each of the enumeration constants, and the object *x*, have *essential type* of *enum<i>*:

```
enum ETAG { A = 8, B = 64, C = 128 } x;
```

The following *anonymous enum type* defines a set of constants. The *essential type* of each constant is the *STLR*. Therefore on a machine with 8-bit *char* and 16-bit *short* types, *A* and *B* have an *essential type* of *signed char* but *C* has an *essential type* of *signed short*.

```
enum { A = 8, B = 64, C = 128 };
```

D.6.3 Character constants

The *standard type* of a character constant (e.g. 'a', 'xy') is *int*, not *char*.

1. If a character constant consists of a single character then its *essential type* is *char*;
2. Otherwise, the *essential type* is the same as its *standard type*.

D.6.4 Boolean constants

The C Standard provides no syntax to explicitly define a constant of *_Bool* type. The library header `<stdbool.h>` defines *false* and *true* in terms of constants 0 and 1 of type *int*, but these have an *essential type* of *essentially Boolean*. If these are not used or are not available, a *constant expression* of type *_Bool* may be declared using an explicit cast or an operator that delivers an *essentially Boolean* value.

Tools may also provide additional ways of identifying *essentially Boolean types*. For example, a tool could be configured to recognize:

```
enum Bool {False, True};
```

The use of a cast to define a *constant expression* of Boolean type (e.g. `(_Bool) 0`) is only appropriate if the expression is not to be used in a `#if` or `#elif` preprocessing directive. Use of a cast within a preprocessing directive is a syntax error.

D.7 The essential type of expressions

The *essential type* of any expression not listed in this section is the same as its *standard type*.

D.7.1 Parenthesis (())

The *essential type* of the result is the *essential type* of the operand.

D.7.2 Comma (,)

The *essential type* of the result is the *essential type* of the right hand operand.

D.7.3 Relational (< <= >= >), Equality (== !=) and Logical (&& || !)

The result of the expression is *essentially Boolean*.

D.7.4 Shift (<< >>)

1. If the left hand operand is *essentially unsigned* then:
 - 1.1 If both operands are *integer constant expressions* then the *essential type* of the result is the *UTLR* of the result;
 - 1.2 Else the *essential type* of the result is the *essential type* of the left hand operand.
2. Else the *essential type* is the *standard type*.

D.7.5 Bitwise complement (~)

1. If the operand is *essentially unsigned* then:
 - 1.1 If the operand is an *integer constant expression* then the *essential type* of the result is the *UTLR* of the result;
 - 1.2 Else the *essential type* of the result is the *essential type* of the operand.
2. Else the *essential type* is the *standard type*.

D.7.6 Unary plus (+)

1. If the operand is *essentially signed* or *essentially unsigned* then the *essential type* of the result is the *essential type* of the operand;
2. Else the *essential type* is the *standard type*.

D.7.7 Unary minus (-)

1. If the operand is *essentially signed* then:
 - 1.1 If the expression is an *integer constant expression* then the *essential type* of the result is the *STLR* of the whole expression;
 - 1.2 Else the *essential type* of the result is the *essential type* of the operand.
2. Else the *essential type* is the *standard type*.

D.7.8 Conditional (?:)

1. If the *essential type* of the second and third operands is the same then the result has the same *essential type*;
2. Else if the second and third operands are both *essentially signed* then the *essential type* of the result is the *essential type* of the one with the highest rank;
3. Else if the second and third operands are both *essentially unsigned* then the *essential type* of the result is the *essential type* of the one with the highest rank;
4. Else the *essential type* is the *standard type*.

D.7.9 Operations subject to the usual arithmetic conversions (* / % + - & | ^)

1. If the operator is + and one operand is *essentially character* and the other is *essentially signed* or *essentially unsigned* having a rank lower than or equal to that of *int* then the *essential type* of the result is *char*;
2. Else if the operator is - and the first operand is *essentially character* type and the second is *essentially signed* or *essentially unsigned* having a rank lower than or equal to that of *int* then the *essential type* of the result is *char*;
3. Else if the operands are both *essentially signed* then:
 - 3.1 If the expression is an *integer constant expression* then the *essential type* of the result is the *STLR* of the result;
 - 3.2 Else the *essential type* of the result is the *essential type* of the operand with the highest rank.
4. Else if the operands are both *essentially unsigned* then:
 - 4.1 If the expression is an *integer constant expression* then the *essential type* of the result is the *UTLR* of the result;
 - 4.2 Else the *essential type* of the result is the *essential type* of the operand with the highest rank.
5. Else the *essential type* is the *standard type*.

D.7.10 Generic selection (`_Generic`)

The *essential type* of a `_Generic` selection is the *essential type* of the selected expression.

Appendix E Applicability to automatically generated code

E.1 Guideline categories for automatically generated code

Most MISRA C guidelines have the same category for automatically generated code as they do for manually generated code.

This section lists those guidelines for which the category is different. For convenience, guidelines with similar rationales are grouped together.

E.1.1 Additional categories

An additional guideline category is used in automatically generated code:

E.1.1.1 Readability

If a guideline has the “readability” category then it is not necessary to comply with it **provided that** the automatically generated code is not intended to be read, reviewed or modified by human programmers. If the code **is** being used by humans then the category remains the same as for manually generated code.

E.1.2 Hiding identifiers

Rule 5.3	Advisory
----------	----------

An automatic code generator is capable of tracking the identifiers that it uses and should not be confused about any identifier reuse. The guideline is therefore advisory when applied to automatically generated code.

When manually generated code is injected into automatically generated code or *vice-versa* it is possible for an identifier to be hidden without the code generator being aware. This guideline therefore remains required if the identifier that is being hidden, or that is hiding another identifier, is declared in manually generated code.

E.1.3 Octal constants

Rule 7.1	Advisory
----------	----------

An automatic code generator is unlikely to generate an octal constant unintentionally.

E.1.4 Compatible declarations with external linkage

Rule 8.4	Advisory	Rule 8.5	Advisory
----------	----------	----------	----------

These guidelines require:

- A compatible declaration to be visible when an object or function with external linkage is defined;
- Each object or function with external linkage to be declared in only one file.

Together, they provide a mechanism for ensuring that all translation units using a declaration of the object or function are doing so in a manner that is compatible with the definition. However, this is not the only strategy for guaranteeing compatibility. For example, an automatic code generator might hold the declaration of each object or function in a data dictionary and use that declaration in each translation unit that needs it.

E.1.4.1 The restrict type qualifier

Rule 8.14	Advisory
-----------	----------

When an automatic code generator can determine that two pointers do not alias it may wish to use the *restrict* type qualifier.

E.1.5 Essential type

Rule 10.1	Advisory	Rule 10.2	Advisory	Rule 10.3	Advisory	Rule 10.4	Advisory
Rule 10.6	Advisory	Rule 10.7	Advisory	Rule 10.8	Advisory	Rule 14.4	Advisory
Rule 20.8	Advisory						

The *essential type* concept is one method for avoiding problems associated with *integer promotion* and the *usual arithmetic conversions*. An automatic code generator may adopt a different approach to avoiding such problems which is equally valid. Further, the automatic code generator is assumed to be aware of both C's implicit type conversion rules and the sizes of the types on the target machine.

The rules relating to *essential type* are therefore all given advisory category when applied to automatically generated code.

E.1.6 Loop counters

Rule 14.1	Advisory
-----------	----------

An automatic code generator is aware of the floating-point implementation and will therefore be aware of the implications of using a floating-point *loop counter*.

E.1.7 Labels and goto

Rule 15.2	Advisory	Rule 15.3	Advisory
-----------	----------	-----------	----------

An automatic code generator may need to generate a backwards jump, for example in the implementation of a state-machine. It is very unlikely to generate a jump into a block of injected manually-generated code and should therefore be able to avoid jumping over initialization.

E.1.8 Switch statements

Rule 16.1	Advisory	Rule 16.2	Advisory	Rule 16.3	Advisory	Rule 16.4	Advisory
Rule 16.5	Advisory	Rule 16.6	Advisory	Rule 16.7	Advisory		

If an automatic code generator determines that it can optimize away a *default clause* that is suggested by the model, for example every value for the controlling expression is covered by a *case clause*, then it may only do so if it inserts a comment into the generated C that explains why the *default clause* is absent. This comment can be reviewed and accepted as part of the MISRA C compliance argument.

An automatic code generator is capable of handling *switch clauses* that fall through into subsequent clauses.

Unusual *switch* statements such as those that have a controlling expression with *essentially Boolean* type, or that have only one *switch clause*, might be indicative of an error in human generated code but may be necessary or desirable in automatically generated code.

E.1.9 Readability

Dir 4.5	Readability	Rule 2.3	Readability	Rule 2.4	Readability	Rule 2.5	Readability
Rule 2.6	Readability	Rule 2.7	Readability	Rule 5.9	Readability	Rule 7.2	Readability
Rule 7.3	Readability	Rule 9.2	Readability	Rule 9.3	Readability	Rule 9.5	Readability
Rule 11.9	Readability	Rule 13.3	Readability	Rule 14.2	Readability	Rule 15.7	Readability
Rule 17.5	Readability	Rule 17.7	Readability	Rule 17.8	Readability	Rule 18.5	Readability
Rule 20.5	Readability						

These guidelines are for the benefit of human developers and reviewers.

If any identifier with internal linkage is declared in injected code, then Rule 5.9 as applied to that identifier is advisory for the entire translation unit.

If code is injected into an automatically generated *for* loop then Rule 14.2 is required insofar as the injected code causes a non-compliance.

E.2 Documentation requirements for automatic code generation tools

The developer of an automatic code generation tool shall provide documentation for each item listed in this section.

E.2.1 Implementation-defined behaviour and language extensions

In accordance with Dir 1.1, the use of any implementation-defined behaviour on which automatically generated code depends shall be documented.

In accordance with Rule 1.2, the use of language extensions by automatically generated code, and the methods by which their use is assured, shall be documented.

In accordance with Dir 4.2, the use of assembly language by automatically generated code shall be documented.

E.2.2 The essential type model

If an automatic code generator does not use the *essential type* model, then the strategy that it uses in its place shall be documented.

E.2.3 Run-time errors

An automatic code generator should select an appropriate code generation strategy to minimize the possibility for run-time failures, as required by Dir 4.1. If the possibility for run-time failures remains, the error handling strategy used by the code generator shall be documented. Since it is likely that manually generated code will be needed in order to handle the error, the error handling interface shall be documented as part of the overall error handling strategy.

Appendix F Obsolescent language features

Obsolescent language features are those identified in the *Future language directions* and *Future library directions* sections of the applicable C Standard — this appendix should be used in conjunction with Rule 1.5.

In the following table:

- **ID** is a unique MISRA sequential identifier
- **Obsolescent Feature** identifies the language feature that has been declared as obsolescent.
- An X in the **C99**, **C11** or **C18** column denotes that the specified feature has been declared as obsolescent in the applicable edition of the C Standard.
- **Comments** identifies other MISRA C guidelines for which the use of the specified feature should also be considered a violation, or specifies a note.

ID	Obsolescent Feature	C99	C11	C18	Comments
1	Declaring an identifier with internal linkage at file scope without the <code>static</code> storage class specifier is an obsolescent feature.	X	X	X	Rule 8.8
2	Restriction of the significance of an external name to fewer than 255 characters [...] is an obsolescent feature.	X	X	X	Note 1
3	The placement of a storage-class specifier other than at the beginning of the declaration specifiers in a declaration is an obsolescent feature.	X	X	X	-
4	The use of function declarators with empty parentheses (not prototype-format parameter type declarators) is an obsolescent feature.	X	X	X	Rule 8.2
5	The use of function definitions with separate parameter identifier and declaration lists (not prototype-format parameter type and identifier declarators) is an obsolescent feature.	X	X	X	Rule 8.2
6	The macro <code>ATOMIC_VAR_INIT</code> is an obsolescent feature.	-	-	X	Rule 1.4
7	The ability to undefine and perhaps then redefine the macros <code>bool</code> , <code>true</code> , and <code>false</code> is an obsolescent feature.	X	X	X	Rule 21.1
8	The <code>gets</code> function is obsolescent, and is deprecated.	X	-	-	Rule 21.6, Note 2
9	The use of <code>ungetc</code> on a binary stream where the file position indicator is zero prior to the call is an obsolescent feature.	X	X	X	Rule 21.6
10	Invoking <code>realloc</code> with a size argument equal to zero is an obsolescent feature.	-	-	X	Rule 21.3

Notes:

1. This is a feature of the implementation and is not (usually) determinable by a static analysis tool.
2. The `gets` function was removed from the C Standard in C11.

Appendix G Implementation-defined behaviour checklist

This Appendix provides the checklist of implementation-defined behaviours referred to in Dir 1.1.

G.1 C90

The numbers in the first column of the table identify the number of an item in the relevant section of Annex G of the C90 Standard [19]. The numbering starts from the beginning of that section and does not include subsequent Technical Corrigenda. So, for example, G.3.1 refers to the first item in section G.3.

C90 Annex	Item	Implementation-defined behaviour
G.3.1	1	How a diagnostic is identified (5.1.1.3).
G.3.2	1	The semantics of the arguments to <code>main</code> (5.1.2.2.1).
G.3.3	1	The number of significant initial characters (beyond 31) in an identifier without external linkage (6.1.2).
	2	The number of significant initial characters (beyond 6) in an identifier with external linkage (6.1.2).
	3	Whether case distinctions are significant in an identifier with external linkage (6.1.2).
G.3.4	1	The members of the source and execution character sets, except as explicitly specified in this International Standard (5.2.1).
	3	The number of bits in a character in the execution character set (5.2.4.2.1).
	4	The mapping of members of the source character set (in character constants and string literals) to members of the execution character set (6.1.3.4).
	5	The value of an integer character constant that contains a character or escape sequence not represented in the basic execution character set or the extended character set for a wide character constant (6.1.3.4).
	6	The value of an integer character constant that contains more than one character or a wide character constant that contains more than one multibyte character (6.1.3.4).
	7	The current locale used to convert multibyte characters into corresponding wide characters (codes) for a wide character constant (6.1.3.4).
G.3.5	1	The representations and sets of values of the various types of integers (6.1.2.5).
	4	The sign of the remainder on integer division (6.3.5).
G.3.6	1	The representations and sets of values of the various types of floating-point numbers (6.1.2.5).
	2	The direction of truncation when an integral number is converted to a floating-point number that cannot exactly represent the original value (6.2.1.3).
	3	The direction of truncation or rounding when a floating-point number is converted to a narrower floating-point number (6.2.1.4).
	4	The order of allocation of bit-fields within a unit (6.5.2.1).
	5	Whether a bit-field can straddle a storage-unit boundary (6.5.2.1).
G.3.10	1	What constitutes an access to an object that has volatile-qualified type (6.5.5.3).
G.3.13	1	Whether the value of a single-character character constant in a constant expression that controls conditional inclusion matches the value of the same character constant in the execution character set. Whether such a character constant may have a negative value (6.8.1).
	2	The method for locating includable source files (6.8.2).

C90 Annex	Item	Implementation-defined behaviour
	3	The support of quoted names for includable source files (6.8.2).
	4	The mapping of source file character sequences (6.8.2).
	5	The behavior on each recognized <code>#pragma</code> directive (6.8.6).
	6	The definitions for <code>__DATE__</code> and <code>__TIME__</code> when respectively, the date and time of translation are not available (6.8.8).
G.3.14	5	Whether the mathematics functions set the integer expression <code>errno</code> to the value of the macro <code>ERANGE</code> on underflow range errors (7.5.1).
	30	The set of environment names and the method for altering the environment list used by the <code>getenv</code> function (7.10.4.4).

G.2 C99 and C11

The numbers in the first column of the table identify the number of an item in the relevant section of Annex J of the C99 Standard [23] or C11 Standard [27], as appropriate. The numbering starts from the beginning of that section and does not include subsequent Technical Corrigenda. So, for example, J.3.1 refers to the first item in section J.3.

Annex	Annex Item		Implementation-defined behaviour
	C99	C11	
J.3.1	1	1	How a diagnostic is identified.
J.3.2	2	2	The name and type of the function called at program startup in a freestanding environment.
	3	3	The effect of program termination in a freestanding environment.
	4	4	An alternative manner in which the <code>main</code> function may be defined.
	5	5	The values given to the strings pointed to by the <code>argv</code> argument to <code>main</code> .
	10	11	The set of environment names and the method for altering the environment list used by the <code>getenv</code> function.
	11	12	The manner of execution of the string by the <code>system</code> function.
J.3.3	1	1	Which additional multibyte characters may appear in identifiers and their correspondence to universal character names.
	2	2	The number of significant initial characters in an identifier.
J.3.4	1	1	The number of bits in a byte.
	2	2	The values of the members of the execution character set.
	4	4	The value of a <code>char</code> object into which has been stored any character other than a member of the basic execution character set.
	6	6	The mapping of members of the source character set (in character constants and string literals) to members of the execution character set.
	7	7	The value of an integer character constant containing more than one character or containing a character or escape sequence that does not map to a single-byte execution character.
	8	8	The value of a wide character constant containing more than one multibyte character, or containing a multibyte character or escape sequence not represented in the extended execution character set.
	9	9	The current locale used to convert a wide character constant consisting of a single multibyte character that maps to a member of the extended execution character set into a corresponding wide character code.
	10	11	The current locale used to convert a wide string literal into corresponding wide character codes.

Annex	Annex Item		Implementation-defined behaviour
	C99	C11	
	11	12	The value of a string literal containing a multibyte character or escape sequence not represented in the execution character set.
J.3.5	1	1	Any extended integer types that exist in the implementation.
	2	2	Whether signed integer types are represented using sign and magnitude, two's complement, or ones' complement, and whether the extraordinary value is a trap representation or an ordinary value.
J.3.6	1	1	The accuracy of the floating-point operations and of the library functions in <code><math.h></code> and <code><complex.h></code> that return floating-point results.
	2	3	The rounding behaviors characterized by non-standard values of <code>FLT_ROUNDS</code> .
	3	4	The evaluation methods characterized by non-standard negative values of <code>FLT_EVAL_METHOD</code> .
	4	5	The direction of rounding when an integer is converted to a floating-point number that cannot exactly represent the original value.
	5	6	The direction of rounding when a floating-point number is converted to a narrower floating-point number.
	6	7	How the nearest representable value or the larger or smaller representable value immediately adjacent to the nearest representable value is chosen for certain floating constants.
	7	8	Whether and how floating expressions are contracted when not disallowed by the <code>FP_CONTRACT</code> pragma.
	8	9	The default state for the <code>FENV_ACCESS</code> pragma.
	9	10	Additional floating-point exceptions, rounding modes, environments, and classifications, and their macro names.
	10	11	The default state for the <code>FP_CONTRACT</code> pragma.
	11		Whether the "inexact" floating-point exception can be raised when the rounded result actually does equal the mathematical result in an IEC 60559 conformant implementation.
	12		Whether the "underflow" (and "inexact") floating-point exception can be raised when a result is tiny but not inexact in an IEC 60559 conformant implementation.
J.3.9	2	2	Allowable bit-field types other than <code>_Bool</code> , <code>signed int</code> , and <code>unsigned int</code> .
	3	4	Whether a bit-field can straddle a storage-unit boundary.
	4	5	The order of allocation of bit-fields within a unit.
J.3.10	1	1	What constitutes an access to an object that has volatile-qualified type.
J.3.11	*	1	The locations within <code>#pragma</code> directives where header name preprocessing tokens are recognized.
	1	2	How sequences in both forms of header names are mapped to headers or external source file names.
	2	3	Whether the value of a character constant in a constant expression that controls conditional inclusion matches the value of the same character constant in the execution character set.
	3	4	Whether the value of a single-character character constant in a constant expression that controls conditional inclusion may have a negative value.
	4	5	The places that are searched for an included <code>< ></code> delimited header, and how the places are specified or the header is identified.
	5	6	How the named source file is searched for in an included <code>" "</code> delimited header.
	6	7	The method by which preprocessing tokens (possibly resulting from macro expansion) in a <code>#include</code> directive are combined into a header name.

Annex	Annex Item		Implementation-defined behaviour
	C99	C11	
	8	9	Whether the # operator inserts a \ character before the \ character that begins a universal character name in a character constant or string literal.
	9	10	The behavior on each recognized non-STDC #pragma directive.
	10	11	The definitions for __DATE__ and __TIME__ when respectively, the date and time of translation are not available.
J.3.12	1	1	Any library facilities available to a freestanding program, other than the minimal set required by clause 4.
	4	4	Whether the <code>feraiseexcept</code> function raises the “inexact” floating-point exception in addition to the “overflow” or “underflow” floating-point exception.
	5	5	Strings other than "C" and "" that may be passed as the second argument to the <code>setlocale</code> function.
	6	6	The types defined for <code>float_t</code> and <code>double_t</code> when the value of the <code>FLT_EVAL_METHOD</code> macro is less than 0.
	9	9	The values returned by the mathematics functions on underflow range errors, whether <code>errno</code> is set to the value of the macro <code>ERANGE</code> when the integer expression <code>math_errhandling & MATH_ERRNO</code> is nonzero, and whether the “underflow” floating-point exception is raised when the integer expression <code>math_errhandling & MATH_ERREXCEPT</code> is nonzero.
	11	12	The base-2 logarithm of the modulus used by the <code>remquo</code> functions in reducing the quotient.
	33	35	The meaning of any n-char or n-wchar sequence in a string representing a NaN that is converted by the <code>strtod</code> , <code>strtof</code> , <code>strtold</code> , <code>wctod</code> , <code>wctof</code> , or <code>wctold</code> function.
	34	36	Whether or not the <code>strtod</code> , <code>strtof</code> , <code>strtold</code> , <code>wctod</code> , <code>wctof</code> , or <code>wctold</code> function sets <code>errno</code> to <code>ERANGE</code> when underflow occurs.
	36	38	Whether open streams with unwritten buffered data are flushed, open streams are closed, or temporary files are removed when the <code>abort</code> or <code>_Exit</code> function is called.
	37	39	The termination status returned to the host environment by the <code>abort</code> , <code>exit</code> , or <code>_Exit</code> function. Or <code>quick_exit</code> for C11.
44	46	Whether the functions in <code><math.h></code> honor the rounding direction mode in an IEC 60559 conformant implementation, unless explicitly specified otherwise.	
J.3.13	1	1	The values or expressions assigned to the macros specified in the headers <code><float.h></code> , <code><limits.h></code> , and <code><stdint.h></code> .
	2	3	The number, order, and encoding of bytes in any object (when not explicitly specified in this International Standard).

Appendix H Undefined and critical unspecified behaviour

This Appendix identifies the undefined and critical unspecified behaviours that are referred to by Rule 1.3.

H.1 Undefined behaviour

The columns in this table have the following meanings:

- "C90 Id" is the number of the undefined behaviour in Annex G of the C90 Standard. The numbering starts from the beginning of section G.2 and does not include subsequent Technical Corrigenda. Where the behaviour is mentioned in the body of the C Standard but not listed in Annex G, a * character is shown;
- "C99 Id" is the number of the undefined behaviour in Annex J of the C99 Standard. The numbering starts from the beginning of section J.2 and does not include subsequent Technical Corrigenda;
- "C11 Id" is the number of the undefined behaviour in Annex J of the C11 Standard. The numbering starts from the beginning of section J.2 and does not include subsequent Technical Corrigenda;;
- "Decidable?" is "Yes" or "No" according to whether detecting instances of the behaviour is, in general, decidable or not;
- "Guidelines" lists the MISRA C Guidelines which, if complied with, avoid the undefined behaviour;
- "Notes" provides additional notes on the behaviour including information on rules that may help to avoid the behaviour even if it cannot be totally avoided.

If a particular undefined behaviour has no entry in the "Guidelines" column then an instance of that behaviour in a program is a violation of Rule 1.3.

Note: it is assumed that any code that is written in other language and linked with the program does not introduce undefined behaviour directly or indirectly. For example, assembly language modules might define overlapping objects which, if accessed from C, would lead to undefined behaviour even though this might not be apparent from the C source code.

Id			Decidable	Guidelines	Notes
C90	C99	C11			
	1	1		N/A	This behaviour is listed in C99 but each such instance is also given its own entry in Annex J. The entry for this behaviour is therefore redundant.
1	2	2	Yes		
2			Yes		
3			Yes	Rule 20.10	
	3	3	Yes		
	4	4	Yes		
		5	No	Dir 5.1, Rule 9.7	
	5	6	Yes		
	6	7	Yes		

Id			Decidable	Guidelines	Notes
C90	C99	C11			
5			Yes	Rule 5.2	
6			Yes	Rule 17.3	
8	7	8	Yes		
9	8	9	No	Dir 4.12, Rule 18.6, Rule 18.9, Rule 21.3, Rule 22.13, Rule 22.14, Rule 22.15, Rule 22.20	
	9	10	No	Dir 4.12, Rule 18.6, Rule 21.3, Rule 22.15	
	10	11	No	Rule 22.13	Compliance with Rule 9.1 also avoids a common cause of this undefined behaviour but it is not sufficient to avoid all situations in which an indeterminate value might arise.
	11	12	No		The following rules help to avoid this behaviour: Rule 9.1, Rule 11.2, Rule 11.3, Rule 11.4, Rule 11.5 and Rule 19.1. However, if a trap representation is copied into an object that does not have character type, for example using memmove, memcpy or via a pointer to character type as permitted by the exception of Rule 11.3, it is not possible to avoid this behaviour.
	12	13	No	Rule 11.2, Rule 11.3, Rule 11.4, Rule 11.5	
	13	14	No		The following rules help to avoid this behaviour: Rule 9.1, Rule 10.1, Rule 11.2, Rule 11.3, Rule 11.4, Rule 11.5, and Rule 19.1. However, if the Exception of Rule 11.3 is used then it is not possible to prevent generation of a negative zero.
10	14	15	Yes	Rule 5.6, Rule 5.7, Rule 8.3	
15			No	Dir 4.1, Dir 4.14, Rule 10.3	
		16	No	Rule 18.8	
	15	17	No	Dir 4.1, Dir 4.14, Rule 10.3	
	16	18	No	Dir 4.1, Dir 4.14, Rule 10.3	
	17	19	No	Rule 9.1, Rule 11.2, Rule 11.3, Rule 11.4, Rule 11.5, Rule 19.1	
16	18	20	Yes		
		21	No		
	19	22	Yes		
17	20	23	Yes		
*	21	24	No	Rule 11.1, Rule 11.2, Rule 11.4, Rule 11.6	
	22	25	No	Rule 11.2, Rule 11.3, Rule 11.5	

Id			Decidable	Guidelines	Notes
C90	C99	C11			
27	23	26	No	Rule 11.1	
4	24	27	Yes		
*	25	28	Yes		
	26	29	Yes		
	27	30	Yes		
7	28	31	Yes	Rule 5.1, Rule 5.2, Rule 5.3, Rule 5.4, Rule 5.5	
	29	32	Yes	Rule 21.2	
11			Yes		
12	30	33	No	Rule 7.4, Rule 11.4, Rule 11.8	
13			Yes		
14			Yes	Rule 20.2	
	31	34	Yes	Rule 20.2	
18	32	35	No	Rule 13.2, Rule 13.3, Rule 13.4	
19	33	36	No	Dir 4.1, Dir 4.14	
20			No	Rule 11.3, Rule 11.4, Rule 11.5	
	34	37	No	Rule 11.3, Rule 11.4, Rule 11.5	
	35		Yes	Rule 18.9	
21			Yes		
22	36	38	No	Rule 8.2, Rule 17.3	Rule 17.3 is only applicable to, and only required for, C90.
23			No	Rule 8.2, Rule 17.3	
24			No	Rule 5.6, Rule 5.7, Rule 8.3, Rule 8.4, Rule 8.5, Rule 11.1, Rule 21.2	
25			No	Rule 8.4, Rule 8.5, Rule 11.1, Rule 21.2, Rule 17.3	
	37	39	No	Rule 8.4, Rule 8.5, Rule 11.1, Rule 21.2	
	38	40	No	Rule 8.2	
	39	41	No	Rule 5.6, Rule 5.7, Rule 8.2, Rule 8.3, Rule 8.4, Rule 8.5, Rule 11.1, Rule 21.2	
		42	Yes	Rule 12.6	
26	40	43	No	Dir 4.1, Dir 4.14	
28			Yes	Rule 11.1	
29	41	44	Yes	Rule 11.1, Rule 11.2, Rule 11.6, Rule 11.7	
	42	45	No	Dir 4.1	

Id			Decidable	Guidelines	Notes
C90	C99	C11			
		*	No		Added by C18
30	43	46	No	Dir 4.14, Rule 18.1	
*	44	47	No	Dir 4.14, Rule 18.1	
31	45	48	No	Dir 4.14, Rule 18.2	
	46	49	No	Rule 18.1	
*	47	50	No		
32	48	51	No	Dir 4.14, Rule 10.1, Rule 12.2	
	49	52	No		Compliance with Dir 4.14 avoids some instances of this undefined behaviour that are related to data from external sources. Compliance with Rule 10.1 avoids this undefined behaviour except when the expression being left-shifted has an unsigned type that is promoted to a signed type.
33	50	53	No	Rule 18.3	
34	51	54	No	Rule 19.1	
*	52	55	Yes		
*	53	56	Yes		
*	54	57	Yes		
*	55	58	Yes		
35	56	59	Yes		
36	57	60	Yes		
37	58	61	Yes		
38			Yes	Rule 6.1	
	59	62	No	Rule 18.7	
	60	63	Yes		
39	61	64	No	Rule 11.4, Rule 11.8, Rule 19.2	
40	62	65	No	Rule 11.4, Rule 11.8, Rule 19.2	
41			No	Rule 9.1	
*	63	66	Yes	Rule 17.13	
*	64	67	Yes		
	65	68	No	Rule 8.14	
	66	69	No	Rule 8.14	
	67	70	Yes	Rule 8.10	
		71	No	Rule 17.9	
		72	Yes		
		73	Yes	Rule 8.15	
*	68	74	Yes		
	69	75	No	Rule 18.10	
	70	76	No	Rule 18.10	
	71	77	No	Rule 17.6	
	72	78	Yes		

Id			Decidable	Guidelines	Notes
C90	C99	C11			
*	73	79	Yes	Rule 8.2, Rule 11.1	
*	74	80	No		
*	75	81	Yes		
42			Yes	Rule 9.2	
	76	82	Yes	Rule 9.2	
	77	83	Yes	Rule 9.2	
44	78	84	Yes	Rule 8.6	
	79	85	Yes	Rule 8.2	
*	80	86	Yes		
45	81	87	Yes	Rule 17.1	
43	82	88	Yes	Rule 17.4	
46	83	89	Yes		
		*	Yes		Added by C18
47	84	90	Yes		
48	85	91	Yes	Rule 20.3	
	86	92	Yes		
49			Yes		
50	87	93	Yes	Rule 20.6	
51	88	94	Yes	Rule 20.10	
52	89	95	Yes	Rule 20.10	
53	90	96	Yes		
	91	97	Yes		
	92	98	Yes		
54	93	99	Yes	Rule 21.1	
55	94	100	No		Compliance with Rule 19.1 avoids a common cause of this undefined behaviour but does not prevent copying part of an object to another part of the same object, such as an array.
*	95	101	Yes		
56			Yes	Rule 17.3, Rule 20.1, Rule 20.4, Rule 21.2	
	96	102	Yes	Rule 20.1	
	97	103	Yes	Rule 20.1, Rule 21.2	
	98	104	Yes	Rule 20.4	
57			Yes	Rule 21.1, Rule 21.2	
	99	105	Yes	Rule 21.2	
	100	106	Yes	Rule 21.1, Rule 21.2	
	101	107	Yes	Rule 21.1	
60	102	108	No	Dir 4.11	
*	103	109	No	Dir 4.11, Rule 21.17 Rule 21.18	
61			Yes	Rule 17.3, Rule 21.2	

Id			Decidable	Guidelines	Notes
C90	C99	C11			
62	104	110	Yes		Compliance with Rule 21.1 prevents undefinition of the macro but no rule prevents the macro expansion from being suppressed, e.g. by means of <code>(assert) (E)</code> .
	105	111	Yes		
	106	112	Yes		
63	107	113	No	Dir 4.11, Rule 21.13	
58			Yes	Rule 21.1	
	108	114	Yes		Compliance with Rule 21.1 prevents undefinition of the macro but no rule prevents the macro expansion from being suppressed, e.g. by means of <code>(errno)</code> if it is implemented as a function-like macro. Compliance with Rule 21.2 prevents definition of the identifier <code>errno</code> .
	109	115	No		Compliance with Rule 21.12 avoids some instances of this undefined behaviour but does not prevent floating-point control modes from being changed.
	110	116	No	Rule 21.12	
	111	117	No	Rule 21.12	
	112	118	No	Dir 4.11, Rule 6.3	
90			No	Rule 21.7	
94			No		Compliance with Dir 4.14 avoids some instances of this undefined behaviour that are related to data from external sources.
	113	119	No		Compliance with Dir 4.14 avoids some instances of this undefined behaviour that are related to data from external sources.
*	114	120	No	Rule 21.19	
*	115	121	No	Rule 21.19	
	116	122	Yes	Rule 21.1, Rule 21.2	
	117	123	Yes		
64			Yes	Rule 21.1, Rule 21.2, Rule 21.4	
	118	124	Yes	Rule 21.1, Rule 21.2, Rule 21.4	
65	119	125	Yes	Rule 21.4	
*	120	126	No	Rule 21.4	
66	121	127	No	Rule 21.4	
67			No	Rule 21.4, Rule 21.5	Compliance with either rule is sufficient to avoid the undefined behaviour.
*	122	128	No	Rule 21.5	
*	123	129	No	Rule 21.5	
		130	No	Rule 21.5	
	124	131	No	Rule 21.5	
68			No	Rule 21.5	
	125	132	No	Rule 21.5	
69	126	133	No	Rule 21.5	

Id			Decidable	Guidelines	Notes
C90	C99	C11			
*	127	134	No	Rule 21.5	
		135	Yes		
*	128	136	No		Compliance with Rule 17.1 avoids instances of this undefined behaviour that arise through improper use of the features of <code><stdarg.h></code> .
70	129	137	No	Rule 17.1	
71			Yes	Rule 17.1, Rule 21.1, Rule 21.2	
	130	138	Yes	Rule 17.1, Rule 21.1, Rule 21.2	
75			No	Rule 17.1	
76			No	Rule 17.1	
	131	139	No	Rule 17.1	
	132	140	Yes	Rule 17.1	
73			No	Rule 17.1	
74			No	Rule 17.1	
	133	141	No	Rule 17.1	
	134	142	No	Rule 17.1	
72	135	143	Yes	Rule 17.1	
		*	Yes		Added by C18
59	136	144	Yes		
	137	145	Yes	Rule 7.5	
	138	146	No	Rule 21.6	
	139	147	No	Rule 21.6	
*	140	148	No	Rule 21.6	If Rule 21.6 is deviated then Rule 22.6 provides protection against this undefined behaviour. Rule 21.6 is preferred as it is decidable.
77	141	149	No	Rule 21.6	
	142	150	No	Rule 21.6	
78	143	151	No	Rule 21.6	
*	144	152	No	Rule 21.6	
79			No	Rule 21.6	
85			No	Rule 21.6	
	145	153	No	Rule 21.6	
	146	154	No	Rule 21.6, Rule 21.10	
*	147	155	No	Rule 21.6	
*	148	156	No	Rule 21.6	
83			No	Rule 21.6	
84			No	Rule 21.6	
	149	157	No	Rule 21.6	
82			No	Rule 21.6	
87			No	Rule 21.6	
	150	158	No	Rule 21.6	
*	151	159	No	Rule 21.6	

Id			Decidable	Guidelines	Notes
C90	C99	C11			
	152	160	No	Rule 21.6	
81	153	161	No	Rule 21.6	
97			No	Rule 21.10	
80	154	162	No	Rule 21.6, Rule 21.10	
86	155	163	No	Rule 21.6	
		164	Yes	Rule 21.6	
89	156	165	No	Rule 21.6	
*	157	166	No	Rule 21.6	
	158	167	No	Rule 21.6	
88	159	168	No	Rule 21.6	
*	160	169	No	Rule 21.6	
*	161	170	No	Rule 21.6	
*	162	171	No	Rule 21.6	
*	163	172	No	Rule 21.6	
*	164	173	No	Rule 21.6	
*	165	174	No	Rule 21.6	
*	166	175	No	Rule 21.6	
*	167	176	No	Rule 21.3	
91	168	177	No	Rule 21.3	
		178	Yes	Rule 21.3	Does not apply to C18 as the behaviour is now defined
92	169	179	No	Rule 21.3, Rule 22.2	
*	170	180	No	Rule 21.3	
*	171	181	No	Rule 21.3	
93	172	182	No	Rule 21.8	
	173	183	No	Rule 21.4	
*	174	184	No	Rule 21.19	
		185	No	Rule 21.5, Rule 21.8	
	175	186	No	Rule 21.21	
	176	187	No	Rule 21.9	
	177	188	No	Rule 21.9	
*	178	189	No	Rule 21.9	
95	179	190	No		
96	180	191	No	Dir 4.11, Rule 21.17, Rule 21.18	
	181	192	No	Dir 4.11, Rule 21.18	
*	182	193	No		Compliance with Rule 21.10 avoids this undefined behaviour except in respect of <i>wcsxfrm</i> .
	183	194	No	Dir 4.11	
	184	195	Yes	Rule 21.11, Rule 21.22	
	185	196	Yes	Rule 21.11	
		*	No	Rule 22.18	Added by C18

Id			Decidable	Guidelines	Notes
C90	C99	C11			
		*	No	Rule 21.26	Added by C18
		*	No	Rule 22.16, Rule 22.17, Rule 22.18	Added by C18
		*	No	Rule 22.11	Added by C18
		*	Yes	Rule 22.20	Added by C18
		*	No	Rule 22.12, Rule 22.15, Rule 22.20	Added by C18
		197	No	Rule 21.10	
	186	198	No	Rule 21.6	
	187	199	No	Dir 4.11	
	188	200	No		
	189	201	No	Dir 4.11	
	190	202	No		
	191	203	No		

H.2 Critical unspecified behaviour

The columns in this table have the following meanings:

- "C90 Id" is the number of the unspecified behaviour in Annex G of the C90 Standard; where the behaviour is mentioned in the body of the C Standard but not listed in Annex G, a * character is shown;
- "C99 Id" is the number of the unspecified behaviour in Annex J of the C99 Standard;
- "C11 Id" is the number of the unspecified behaviour in Annex J of the C11 Standard;
- "Critical?" is "Yes" or "No" according to whether reliance on this behaviour is likely to lead to unexpected program operation;
- "Guidelines" lists the MISRA C Guidelines which, if complied with, avoid the unspecified behaviour;
- "Notes" provides additional notes on the behaviour including information on rules that may help to avoid the behaviour even if it cannot be totally avoided.

If a particular unspecified behaviour is marked as being critical but has no entry in the "Guidelines" column then reliance on that behaviour in a program is a violation of Rule 1.3.

Note: it is assumed that any code that is written in other language and linked with the program does not rely on unspecified behaviour directly or indirectly. For example, assembly language functions might attempt to access parameters despite their layout being unspecified.

Id			Critical	Guidelines	Notes
C90	C99	C11			
1	1	1	No		
	2	2	No		
		3	No		

Id			Critical	Guidelines	Notes
C90	C99	C11			
2	3	4	No	Rule 21.6	
3	4	5	No	Rule 21.6	
4	5	6	No	Rule 21.6	
5	6	7	No	Rule 21.6	
6			Yes		
	7	8	Yes	Rule 5.1	
	8	9	Yes		
	9	10	Yes		Compliance with Rule 21.16 avoids this unspecified behaviour in respect of <i>memcmp</i> only.
	10	11	Yes	Rule 19.2	
	11	12	Yes		
	12	13	Yes		
	13	14	Yes		Compliance with Rule 10.1 avoids generation of negative zeros when operating on expressions that have a signed type before promotion.
	14	15	Yes	Rule 7.4	
7,8	15	16	Yes	Rule 13.2	
9	16	17	Yes	Rule 13.2	
	17	18	Yes	Rule 13.1	
7	18	19	Yes	Rule 13.2	
10	19	20	No		
	20	21	Yes	Rule 8.10	
	21	22	Yes	Rule 13.6, Rule 18.8	
7	22	23	Yes	Rule 13.1	
11	23	24	No		
*	24	25	Yes		
12	25	26	Yes	Rule 20.10, Rule 20.11	
13	26		No		
		*	Yes		Added by C18 - <code>#line __LINE__ new-line</code>
	27	27	Yes	Rule 21.12	
	28	28	Yes	Rule 21.12	
	29	29	No		
	30	30	Yes	Dir 4.11, Dir 4.15	
		31	Yes		
	31	32	Yes	Dir 4.11	
		33	Yes		
		34	No		
14	32	35	No	Rule 21.4	
15	33	36	No	Rule 17.1	
	34	37	Yes	Rule 21.6	
16	35	38	Yes	Rule 21.6	
17	36	39	Yes	Rule 21.6	

Id			Critical	Guidelines	Notes
C90	C99	C11			
18	37	40	Yes	Rule 21.6	
	38	41	No		
19	39	42	No	Rule 18.1, Rule 18.2, Rule 18.3, Rule 21.3	Compliance with either Rule 21.3 or all of Rule 18.1, Rule 18.2 and Rule 18.3 will avoid this unspecified behaviour.
	40	43	Yes	Rule 21.3	
		44	Yes		
		45	Yes		
20	41	46	Yes	Rule 21.9	C11 incorrectly omitted <i>align_alloc</i> , which was corrected in C18.
21	42	47	Yes	Rule 21.9	C11 incorrectly omitted <i>align_alloc</i> , which was corrected in C18.
22	43	48	Yes	Rule 21.10	
	44	49	Yes	Rule 21.10	
		50	Yes		
		*	Yes		Added by C18 – <i>thrd_exit</i> destructor invocation ordering
		*	Yes		Added by C18 – <i>tss_delete</i> destructor invocations with multiple threads
	45	51	Yes		
	46	52	Yes	Dir 4.15	
	47	53	Yes	Dir 4.15	
	TC3	54	Yes	Dir 4.11, Dir 4.15	Added to C99 by TC3.
	TC3	55	Yes	Dir 4.11, Dir 4.15	Added to C99 by TC3.
	48	56	Yes	Dir 4.11	
	49	57	Yes	Dir 4.11	
	50	58	Yes	Dir 4.11	

Appendix I Example deviation record

Withdrawn -- superseded by Appendix B of MISRA Compliance:2020 [11].

Appendix J Glossary

Assigned

An expression is *assigned* if it is the subject of an *assignment*.

Assignment

It is sometimes convenient to use the term *assignment* to denote any operation which takes place as if it were by assignment. The operations covered by this term are:

- Assignment by means of one of the assignment operators;
- Passing an argument to a function, in which case the argument is copied as if by assignment to the corresponding parameter;
- Returning an expression from a function, in which case the result is copied as if by assignment to an object with the function's return type;
- Using an expression to initialize all or part of an object, including a compound literal in C99, in which case the expression is copied as if by assignment to the destination.

Code

Code consists of everything within a translation unit that is not excluded by conditional compilation.

Dead code

Any operation whose result does not affect the program's behaviour is *dead code*.

Note: initialization is not the same as an assignment operation and is therefore not a candidate for *dead code*.

Examples of *dead code* include:

- Storing values into non-*volatile* objects that are never subsequently used;
- An expression statement that does not assign a value to an object.

The following are never *dead code*:

- Accessing a *volatile* object;
- Using a language extension.

Since the time at which actions occur is often an important aspect of the behaviour of an embedded program, it follows that code whose effect is to insert a delay is not *dead code*. However, any such code is very likely to be written with reference to timer hardware and will therefore involve *volatile* object accesses in any event. For example:

```
extern const volatile uint16_t MILLISEC_TIMER;

void delay_ms ( uint16_t n )
{
    uint16_t start_time = MILLISEC_TIMER;

    while ( ( MILLISEC_TIMER - start_time ) < n )
    {
        /* wait */
    }
}
```

External identifier

An *external identifier* is an identifier with external linkage and must therefore denote either an object or a function.

Generic function

The term *generic function* is defined by the C Standard to mean a callable entity that can handle operands of different types in appropriately different ways (i.e. “ad-hoc polymorphism”).

The exact details of its implementation are left intentionally unspecified in order to provide maximum latitude to implementations. It is not required that any *generic functions* in the C Standard library are actually implemented using generic selections, although generic selections provide an in-language mechanism for doing so; compiler intrinsic features are another possible mechanism.

Other properties of the entity representing a *generic function* are also unspecified. It is not likely that a *generic function* has a single address that can be taken, for instance.

Generic match

A type *matches* if it is compatible, without undergoing any implicit conversion, with the type of the value of the controlling expression. At most one type listed in the selection may *match*; the generic association specifying that type will be *selected*.

If no type *matches*, the default association will be *selected* if present.

Generic selection

A generic association is *selected* if the type specified is compatible with the type of the value of the controlling expression. If no explicitly-specified type is compatible with the type of the value of the controlling expression, the default association is *selected* if present. Otherwise, if no association is *selected*, the generic selection violates the constraint that exactly one association must be *selected* by the generic selection expression.

The result expression of the *selected* association becomes the result of the complete generic selection expression.

Header file

A *header file* is any file that is the subject of a *#include* directive.

Note: the filename extension is not significant.

Inline function

An *inline function* is one that is declared with the *inline* function specifier.

Loop counter

A *loop counter* is defined in Section 8.14 .

Macro identifier

A *macro identifier* is an identifier that is either a macro name or a macro parameter.

Opaque type

A type whose implementation detail is not made available to its users.

Persistent side effect

A *side effect* is said to be *persistent* at a particular point in execution if it might have an effect on the execution state at that point. All of the following *side effects* are *persistent* at a given point in the program:

- Modifying a file;
- Modifying an object;
- Accessing a *volatile* object.

When a function is called, it may have *side effects*. Modifying a file or accessing a *volatile* object are persistent as viewed by the calling function. However any objects modified by the called function, whose lifetimes have ended by the time it returns, do not affect the caller's execution state. Any *side effects* arising from modifying such objects are **not persistent** as viewed by the caller.

The determination of whether a function has *persistent side effects* takes no consideration of the possible values for parameters or other non-local objects.

For example:

```
uint16_t f ( void )
{
    uint16_t temp = 1;          /* Side effect of modifying temp is not
                               * persistent to the caller */

    return ( temp );
}

x = f ( );
```

Project

A *project* consists of the *code* from the set of translation units used to build a single executable. A typical embedded controller might consist of several logically distinct projects, for example:

- Primary bootloader;
- Secondary bootloader;
- Main application.

Prototype form

A function type is said to be in *prototype form* if it includes a prototype, *i.e.* it specifies a list of parameter types and, optionally, names. For example, the following are all in *prototype form*:

```
void f ( void );          /* Prototype specifying no parameters */
void g ( int16_t, char ); /* Prototype specifying types but no names */
void h ( uint32_t n );    /* Prototype specifying type and name */
```

The following are not in *prototype form*:

```
void f ( );              /* No information about parameters */

void g ( x, y )          /* Function definition "K&R"-style */
int16_t x;
char y;
{
}
```

Standard type

The *standard type* of an expression is its type as determined according to the C Standard.

Type-punning

Type-punning is when an object is referred to using different types.

For example, writing data to one member of a union and reading it from another.

Uniform resource identifier (URI)

A *uniform resource identifier* (URI) is a compact sequence of characters that identifies an abstract or physical resource, as defined by RFC 3986 [41].

Unreachable code

Code is unreachable if, in the absence of hardware failures and ignoring the effects of undefined behaviour, there is no combination of program inputs that can cause it to be executed.

Unreachable code may result from the structure of the program's control flow graph. For example, any code following a *return* statement is unreachable and some operands of logical or conditional operators may be unreachable:

```
uint16_t y;

uint16_t f ( void )
{
    uint16_t x;

    return true ? y : ( y + 1U );    /* y + 1 is never executed */
    x = 10U;                       /* statement is never executed */
}
```

Unreachable code may also result when there is a control flow path to a statement but that path can never be followed, for example:

```
uint16_t x, y;

switch ( 2u * x )
{
case 0:
    y = 10u;
    break;
case 1:
    y = 0u;    /* unreachable: 2 * x can never be odd */
    break;

case 2:
    y = x;
    break;
default:
    break;
}
```

Used

An expression is said to be *used* if it appears in one of the following contexts:

- The operand of an operator, including a parameter of a function call operator;
- The controlling expression of an *if*, *while*, *for*, *do ... while*, or *switch* statement;
- The value returned by a *return* statement;
- An *initializer*.

This term should not be confused with reading the value of an object, for example:

```
uint16_t x = 3;
uint16_t *p = &x;          /* x is used but not read */

return *p;                 /* x is read but not used */
```

Appendix K Change log

Item	Revisions	Reason
Front material	Revised for 3rd Edition, 1st Revision	3rd Ed 1st Rev
	Updated for 3rd Edition, 2nd Revision	3rd Ed 2nd Rev
Section 1	Split into sub-sections “Background” and “The vision”	AMD2
Section 1.1	Added reference to <i>high security</i>	AMD1
	New sub-section title	AMD2
Section 1.2	New sub-section with updated vision	AMD2
Section 1.3	New sub-section, by moving part of former s3.1 and updated to include C11/C18	AMD2
	Revised presentation of supported editions of the C Standard	3rd Ed 2nd Rev
Section 1.4	New sub-section, by moving former s5.1 and updated for MISRA Compliance	AMD2+Compliance
Section 2.2.1	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Section 2.3	Added to explain that MISRA is applicable to security	AMD1
Section 3	Deleted — replaced with placeholder	AMD2+Compliance
Section 4	Deleted — replaced with placeholder	AMD2+Compliance
Section 5	Deleted — replaced with placeholder	AMD2+Compliance
Section 6.2.2	Replaced internal reference with reference to MISRA Compliance	AMD2+Compliance
Section 6.5	Replaced internal reference with reference to MISRA Compliance	AMD2+Compliance
Section 6.7	Terminology: <i>the C Standard, the Standard Library</i>	3rd Ed 2nd Rev
Section 6.9	Made Cxx definition generic	AMD2
	Added note linking C18 to C11	AMD2
	Clarified <i>Applies to</i>	TC2
	Added example complex type definitions	AMD3
	Added explanation of <i>Example</i> and <i>See also</i> sections	AMD4
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Section 6.10.1	Revised introductory text, table headings and C Standard references	AMD2
	Terminology: <i>the C Standard, edition, portability issues</i>	3rd Ed 2nd Rev
Section 6.10.2	Added reference to C Secure to table	AMD1
	Terminology: <i>the C Standard, portability issues</i>	3rd Ed 2nd Rev
Dir 1.1	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
	Editorial: formatting of Rationale sub-headings	3rd Ed 2nd Rev
Dir 2.1	Added C11 references, removed internal link to tool configuration	AMD2
Dir 3.1	Added C11 references	AMD2
Dir 4.1	Added C11 references, removed internal link to process requirements	AMD2
	Editorial: formatting of Rationale sub-headings	3rd Ed 2nd Rev
Dir 4.2	Added C11 references	AMD2
Dir 4.3	Added C11 references, revised mentions of C90/C99	AMD2

Item	Revisions	Reason
Dir 4.4	Added C11 references	AMD2
Dir 4.5	Added C11 references	AMD2
Dir 4.6	Clarification regarding plain <i>char</i>	TC1
	Added C11 references, revised mentions of C90/C99	AMD2
	Added example complex type definitions	AMD3
	Terminology: <i>the Standard Library</i>	3rd Ed 2nd Rev
Dir 4.7	Added C11 references	AMD2
	Added example complex typedefs and supporting narrative	AMD3
	Terminology: <i>the Standard Library</i>	3rd Ed 2nd Rev
Dir 4.8	Clarification regarding multiple pointers	TC1
	Added C11 references	AMD2
Dir 4.9	Added C11 references, revised mentions of C90/C99	AMD2
	Added exception for Generic macros	AMD3
Dir 4.10	Added C11 references	AMD2
	Clarified that other mechanisms are allowed	TC2
Dir 4.11	Correction of C99 reference	TC1
	Added C11 references	AMD2
	Extended rationale for precision	AMD3
	Terminology: <i>the C Standard, the Standard Library</i>	3rd Ed 2nd Rev
Dir 4.12	Added C11 references	AMD2
	Terminology: <i>the C Standard, the Standard Library</i>	3rd Ed 2nd Rev
Dir 4.13	Added C11 references	AMD2
Dir 4.14	New Directive	AMD1
	Added C11 references	AMD2
Dir 4.15	Added new Directive	AMD3
Section 7.5	New section	AMD4
Dir 5.1	New Directive	AMD4
Dir 5.2	New Directive	AMD4
Dir 5.3	New Directive	AMD4
Rule 1.1	Added C11 references, updated internal reference to C Standard definition	AMD2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Rule 1.2	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 1.3	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 1.4	New Rule	AMD2
	Removed restrictions on features subject to new AMD3 guidance	AMD3
	Removed restrictions on features subject to new AMD4 guidance	AMD4
Rule 1.5	New Rule	AMD3
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev

Item	Revisions	Reason
Rule 2.1	Added C11 references	AMD2
Rule 2.2	Clarification regarding initialization	TC1
	Added C11 references	AMD2
	Clarified that a <i>used</i> cast is not <i>dead code</i>	TC2
	Revised headline	AMD4
Rule 2.3	Added C11 references	AMD2
Rule 2.4	Added C11 references	AMD2
Rule 2.5	Clarification regarding <code>#undef</code>	TC1
	Added C11 references	AMD2
	Used <code>define</code> in place of <code>declare</code>	TC2
Rule 2.6	Added C11 references	AMD2
Rule 2.7	Added C11 references	AMD2
	Revised headline	AMD4
Rule 2.8	New Rule	AMD4
Rule 3.1	Added C11 references	AMD2
	Permit URIs in comments	AMD4
Rule 3.2	Added C11 references, revised mentions of C90/C99	AMD2
Rule 4.1	Added C11 references, revised mentions of C90/C99	AMD2
Rule 4.2	Added C11 references	AMD2
Rule 5.1	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Rule 5.2	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Rule 5.3	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Rule 5.4	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Rule 5.5	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Rule 5.6	Added C11 references	AMD2
Rule 5.7	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 5.8	Added C11 references	AMD2
Rule 5.9	Clarification regarding identifier scope	TC1
	Added C11 references	AMD2
Rule 6.1	Added C11 references, revised mentions of C90/C99	AMD2
Rule 6.2	Added C11 references, revised mentions of C90/C99	AMD2
Rule 6.3	New Rule	AMD3
Rule 7.1	Added C11 references	AMD2
Rule 7.2	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 7.3	Added C11 references	AMD2
Rule 7.4	Added C11 references, revised mentions of C90/C99	AMD2

Item	Revisions	Reason
	Exception when using variadic functions	TC2
Rule 7.5	New Rule	AMD3
	Added <i>See also</i> to Rule 7.6	AMD4
Rule 7.6	New Rule	AMD4
Rule 8.1	Added C11 references	AMD2
Rule 8.2	Added C11 references	AMD2
	Added <i>See also</i> to Rule 8.3	TC2
	Added <i>See also</i> to Rule 1.5	AMD3
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Rule 8.3	Added C11 references	AMD2
	Added <i>See also</i> to Rule 8.2 and new <i>Exception</i>	TC2
Rule 8.4	Added an <i>Exception</i> for function <i>main()</i>	TC1
	Added C11 references	AMD2
Rule 8.5	Added C11 references	AMD2
Rule 8.6	Added C11 references	AMD2
	Added <i>See also</i>	AMD4
Rule 8.7	Added C11 references	AMD2
	Added <i>Example</i>	TC2
Rule 8.8	Added C11 references	AMD2
	Added <i>See also</i> to Rule 1.5	AMD3
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 8.9	Added C11 references	AMD2
	Terminology: <i>declared not defined</i>	AMD4
Rule 8.10	Added C11 references	AMD2
Rule 8.11	Added C11 references	AMD2
Rule 8.12	Added C11 references	AMD2
Rule 8.13	Added C11 references	AMD2
Rule 8.14	Added C11 references	AMD2
	Terminology: <i>the Standard Library</i>	3rd Ed 2nd Rev
Rule 8.15	New Rule	AMD3
Rule 8.16	New Rule	AMD3
Rule 8.17	New Rule	AMD3
Rule 9.1	Added C11 references	AMD2
	Added <i>See also</i> to Rule 9.7 and an <i>exception</i> for atomics	AMD4
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 9.2	Added C11 references	AMD2
	Added <i>See also</i> to Rule 9.6	AMD4
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 9.3	Added C11 references	AMD2
Rule 9.4	Added C11 references, revised mentions of C90/C99	AMD2
	Usage of <i>designated initializers</i>	AMD4
Rule 9.5	Added C11 references, revised mentions of C90/C99	AMD2
Rule 9.6	New Rule	AMD4

Item	Revisions	Reason
Rule 9.7	New Rule	AMD4
Section 8.10.1	Clarified applicability to expressions of pointer type	TC2
Section 8.10.2	Revised mentions of C90/C99	AMD2
	Added Complex	AMD3
Rule 10.1	Clarification regarding the use of the logical operators	TC1
	Added C11 references	AMD2
	Clarified use of undefined behaviour	TC2
	Added Complex	AMD3
	Terminology: <i>expressions</i> not <i>objects</i>	AMD4
Rule 10.2	Added C11 references	AMD2
	Clarified use of + and - operators on <i>essentially character types</i>	TC2
	Included complex types	AMD3
Rule 10.3	Correction of C99 reference	TC1
	Clarification regarding <i>switch</i> statement <i>case</i> labels	TC1
	Clarification of Exception 1	TC1
	Added C11 references	AMD2
	Clarified Amplification 2 for <i>switch</i> statement <i>case</i> labels	TC2
	Added Exception for float-to-complex	AMD3
Rule 10.4	Correction of Examples	TC1
	Added C11 references, revised mentions of C90/C99	AMD2
	Added Exception for float-to-complex and complex-to-float	AMD3
Rule 10.5	Clarified Exception with respect to enumerated types	TC1
	Added C11 references, revised mentions of C90/C99	AMD2
	Added Complex	AMD3
Section 8.10.3	Completed list of composite operators	TC1
Rule 10.6	Added C11 references	AMD2
Rule 10.7	Added C11 references	AMD2
	Allowed co-existence of float and Complex	AMD3
Rule 10.8	Corrected typo in Rationale	TC1
	Added C11 references	AMD2
	Allowed casting between float and Complex	AMD3
Section 8.11	Revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 11.1	Added C11 references, revise mentions of C90/C99	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 11.2	Clarification regarding unqualified types that are pointed to by pointers	TC1
	Added C11 references	AMD2
Rule 11.3	Added C11 references, revise mentions of C90/C99	AMD2
	Added cross-reference	TC2
	Use of <i>conversion</i> in place of <i>cast</i> ; use of <i>_Atomic</i>	AMD4
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 11.4	Clarification regarding non-object pointers	TC1
	Added C11 references, revise mentions of C90/C99	AMD2

Item	Revisions	Reason
Rule 11.5	Added C11 references	AMD2
Rule 11.6	Added C11 references	AMD2
	Corrected reference to implementation-defined behaviour	TC2
Rule 11.7	Added C11 references	AMD2
Rule 11.8	Added C11 references, revise mentions of C90/C99	AMD2
	Use of <i>conversion</i> in place of <i>cast</i> ; use of <i>_Atomic</i>	AMD4
Rule 11.9	Added an Exception for the { 0 } initializer	TC1
	Improved the Example comment	TC1
	Added C11 references	AMD2
Rule 11.10	New Rule	AMD4
Rule 12.1	Added C11 references	AMD2
	Added <i>generic</i> and <i>_AlignOf</i>	AMD2
Rule 12.2	Added C11 references	AMD2
Rule 12.3	Added C11 references	AMD2
Rule 12.4	Clarification regarding constant expressions	TC1
	Added C11 references	AMD2
Rule 12.5	New Rule	AMD1
	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 12.6	New Rule	AMD4
Rule 13.1	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
Rule 13.2	Correction of Example	TC1
	Added C11 references, revised mentions of C90/C99	AMD2
	Clarified Amplification and Rationale	TC2
	Terminology: <i>the C Standard, the Standard Library, edition</i>	3rd Ed 2nd Rev
	Extended to address concurrency aspects	AMD4
Rule 13.3	Added C11 references	AMD2
	Terminology: <i>the C Standard</i> ; italicise <i>full expression</i>	3rd Ed 2nd Rev
Rule 13.4	Added C11 references	AMD2
Rule 13.5	Added C11 references	AMD2
Rule 13.6	Added C11 references, revised mentions of C90/C99	AMD2
	Recategorized from Mandatory to Required	TC2
Rule 14.1	Added C11 references	AMD2
Rule 14.2	Clarification regarding loop counter initialization	TC1
	Added C11 references, revised mentions of C90/C99	AMD2
Rule 14.3	Added C11 references	AMD2
	Clarified meaning of <i>false</i> in Exception 2	TC2
Rule 14.4	Added C11 references	AMD2
Rule 15.1	Added C11 references	AMD2
Rule 15.2	Added C11 references	AMD2
Rule 15.3	Added C11 references, revised mentions of C90/C99	AMD2
Rule 15.4	Added C11 references	AMD2

Item	Revisions	Reason
Rule 15.5	Added C11 references	AMD2
Rule 15.6	Correction of Example	TC1
	Added C11 references	AMD2
Rule 15.7	Clarification regarding side effects of function calls	TC1
	Added C11 references	AMD2
	Correct "See also"	TC2
Rule 16.1	Type-setting corrections	TC1
	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 16.2	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 16.3	Added C11 references	AMD2
Rule 16.4	Added C11 references	AMD2
Rule 16.5	Added C11 references	AMD2
Rule 16.6	Added C11 references	AMD2
Rule 16.7	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 17.1	Added C11 references, revised mentions of C90/C99	AMD2
	Revised headline and amplification for consistency of terminology	AMD3
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 17.2	Added C11 references	AMD2
	Terminology: "a serious failure" replaced by "failure"	3rd Ed 2nd Rev
Rule 17.4	Added C11 references, revised mentions of C90/C99	AMD2
	Added an exception for <code>main()</code>	TC2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 17.5	Added C11 references	AMD2
	Clarified categorization as <i>required</i> - use of <i>should</i> replaced	TC2
Rule 17.6	Added C11 references, revised mentions of C90/C99	AMD2
Rule 17.7	Added C11 references	AMD2
Rule 17.8	Added C11 references	AMD2
Rule 17.9	New Rule	AMD3
Rule 17.10	New Rule	AMD3
Rule 17.11	New Rule	AMD3
Rule 17.12	New Rule	AMD3
Rule 17.13	New Rule	AMD3
Rule 18.1	Added C11 references, revised mentions of C90/C99	AMD2
	Clarified interaction with Rule 11.3	TC2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 18.2	Added C11 references	AMD2
Rule 18.3	Added C11 references	AMD2
	Terminology: <i>expressions</i> not <i>objects</i>	AMD4
Rule 18.4	Added C11 references	AMD2
Rule 18.5	Added C11 references	AMD2

Item	Revisions	Reason
	Terminology: "seriously impair" replaced by "impair"	3rd Ed 2nd Rev
Rule 18.6	Added C11 references	AMD2
	Extended to include thread-local storage	AMD4
Rule 18.7	Added C11 references	AMD2
Rule 18.8	Added C11 references	AMD2
	Focused on variable length arrays (and not array types)	AMD4
Rule 18.9	New Rule	AMD3
Rule 18.10	New Rule	AMD4
Rule 19.1	Correction of Example	TC1
	Added C11 references	AMD2
	Terminology: <i>the Standard Library</i>	3rd Ed 2nd Rev
Rule 19.2	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 20.1	Added C11 references	AMD2
	Terminology: <i>the Standard Library</i>	3rd Ed 2nd Rev
Rule 20.2	Added C11 references	AMD2
Rule 20.3	Added C11 references	AMD2
Rule 20.4	Added C11 references, revised mentions of C90/C99	AMD2
Rule 20.5	Added C11 references	AMD2
Rule 20.6	Added C11 references	AMD2
Rule 20.7	Added C11 references	AMD2
Rule 20.8	Added C11 references	AMD2
Rule 20.9	Added C11 references	AMD2
Rule 20.10	Added C11 references	AMD2
Rule 20.11	Added C11 references	AMD2
Rule 20.12	Added C11 references	AMD2
Rule 20.13	Added C11 references	AMD2
Rule 20.14	Added C11 references	AMD2
	Corrected typo in filename	TC2
Rule 21.1	Added cross-reference to See Also	TC1
	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 21.2	Clarification regarding the Rule scope	TC1
	Correction of Example	TC1
	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 21.3	Added C11 references, revised mentions of C90/C99	AMD2
	Terminology: <i>the Standard Library</i>	3rd Ed 2nd Rev
Rule 21.4	Added C11 references	AMD2
	Revised amplification for consistency of terminology	AMD3
Rule 21.5	Added C11 references	AMD2
	Revised amplification for consistency of terminology	AMD3
Rule 21.6	Added C11 references, revised mentions of C90/C99	AMD2

Item	Revisions	Reason
Rule 21.7	Consistency of terminology regarding the Headline	TC1
	Added C11 references, revised mentions of C90/C99	AMD2
Rule 21.8	Consistency of terminology regarding the Headline	TC1
	Removed <i>getenv()</i> from scope	AMD1
	Added cross-references to See Also	AMD1
	Removed <i>system()</i> from scope - see new Rule 21.21	AMD2
	Added <i>_Exit()</i> and <i>quick_exit()</i> to scope	AMD2
Rule 21.9	Consistency of terminology regarding the Headline	TC1
	Added C11 references	AMD2
Rule 21.10	Added C11 references, revised mentions of C90/C99	AMD2
	Revised amplification for consistency of terminology	AMD3
Rule 21.11	Added C11 references	AMD2
	Recategorized as <i>Advisory</i> and linked to new Rules	AMD3
Rule 21.12	Added C11 references	AMD2
	Recategorized as <i>Required</i> and clarified new C11 features	AMD3
Rule 21.13	New Rule	AMD1
	Added C11 references	AMD2
Rule 21.14	New Rule	AMD1
	Added C11 references	AMD2
Rule 21.15	New Rule	AMD1
	Added C11 references	AMD2
Rule 21.16	New Rule	AMD1
	Added C11 references	AMD2
Rule 21.17	New Rule	AMD1
	Added C11 references	AMD2
Rule 21.18	New Rule	AMD1
	Added C11 references	AMD2
Rule 21.19	New Rule	AMD1
	Added C11 references	AMD2
	Removed redundant <i>See also</i> reference	TC2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 21.20	New Rule	AMD1
	Added C11 references	AMD2
	Clarified pairings of functions	TC2
Rule 21.21	New Rule (split from Rule 21.8)	AMD2
Rule 21.22	New Rule	AMD3
Rule 21.23	New Rule	AMD3
Rule 21.24	New Rule	AMD3
Rule 21.25	New Rule	AMD4
Rule 21.26	New Rule	AMD4
Rule 22.1	Added C11 references	AMD2
	Added <i>aligned_alloc()</i> to scope	AMD2
Rule 22.2	Added C11 references	AMD2

Item	Revisions	Reason
Rule 22.3	Added C11 references	AMD2
	Terminology: <i>the C Standard, the Standard Library</i>	3rd Ed 2nd Rev
Rule 22.4	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 22.5	Added C11 references, revised mentions of C90/C99	AMD2
Rule 22.6	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 22.7	New Rule	AMD1
	Added C11 references	AMD2
Rule 22.8	New Rule	AMD1
	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 22.9	New Rule	AMD1
	Added C11 references	AMD2
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 22.10	New Rule	AMD1
	Added C11 references	AMD2
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Rule 22.11	New Rule	AMD4
Section 8.23	New section	AMD3
Rule 23.1	New Rule	AMD3
Rule 23.2	New Rule	AMD3
Rule 23.3	New Rule	AMD3
Rule 23.4	New Rule	AMD3
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Rule 23.5	New Rule	AMD3
Rule 23.6	New Rule	AMD3
Rule 23.7	New Rule	AMD3
Rule 23.8	New Rule	AMD3
Section 9	Added new reference for AMD1 and C-Secure	AMD1
	Added new reference for TC1	TC1
	Added new reference for Compliance:2016	Compliance
	Added new references for AMD2, Compliance:2020 and C11/C18	AMD2
	Added new reference for TC2	TC2

Item	Revisions	Reason
	Added new reference for AMD3; replace IEEE 754 with IEC 60559	AMD3
	Added new reference for AMD4; add AUTOSAR, ARINC 653, OSEK and RFC 3986	AMD4
	Reordered references and group under sub-headings	3rd Ed 2nd Rev
	Updated editions of revised ISO and IEC standards	3rd Ed 2nd Rev
Appendix A	Updated to include new and revised Directives and Rules	3rd Ed 1st Rev
	Updated to include new and revised Directives and Rules	3rd Ed 2nd Rev
Appendix B	Updated to include new and revised Directives and Rules	3rd Ed 1st Rev
	Updated to include new and revised Directives and Rules	3rd Ed 2nd Rev
Appendix C	Added <i>Note</i> regarding implicit and explicit <i>conversions</i>	TC2
Appendix C.1.1	Revised mentions of C90/C99	AMD2
Appendix C.1.3	Added <i>loss of imaginary part</i> and conversion of float to complex	AMD3
Appendix C.2.4	Terminology: <i>hazardous</i> replaced with <i>undesirable</i>	3rd Ed 2nd Rev
Appendix D	Added <i>complex floating-point expressions</i>	AMD3
Appendix D.1	Revised mentions of C90/C99	AMD2
	Clarified the <i>essential type</i> of additional C types	TC2
	Added Complex	AMD3
Appendix D.3	Clarification regarding use of <i>STLR</i> and <i>UTLR</i>	TC1
Appendix D.6	Revised mentions of C90/C99	AMD2
Appendix D.6.4	Revised mentions of C90/C99	AMD2
Appendix D.7	Completed list of operators	TC1
	Numbered sub-paragraphs (were bullets)	3rd Ed 1st Rev
Appendix D.7.9		
Appendix D.7.10	Combined three paragraphs for clarity	TC2
Appendix D.7.11		
Appendix D.7.12 (now D.7.10)	Revised mentions of C90/C99, add <i>_Generic</i>	AMD2
	Renumbered as D.7.9 and clarified	AMD3
Appendix F	Deleted — replaced with placeholder	AMD2+Compliance
	New appendix <i>Obsolescent language features</i>	AMD3
	Terminology: <i>the C Standard, edition</i>	3rd Ed 2nd Rev
Appendix G	Correction of numbering	TC1
	Replaced to add C11 references	AMD2
Appendix H	Replaced to add C11 references	AMD2
	Added new Guidelines and addressed behaviours	AMD3
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev
Appendix I	Deleted — replaced with placeholder	AMD2+Compliance
Appendix J	Clarification regarding persistent side effects	TC1
	Added new <i>generics</i> -related definitions, and <i>type-punning</i>	AMD3
	Added new definition of <i>Uniform Resource Indicator (URI)</i>	AMD4
	Terminology: <i>the C Standard</i>	3rd Ed 2nd Rev

Item	Revisions	Reason
Appendix K	Change Log created	3rd Ed 1st Rev
	Change Log updated	3rd Ed 2nd Rev