

ISO/IEC JTC 1/SC 22/WG 23 N 0276

Revised draft language-specific annex for C

Date	10 September 2010
Contributed by	John Benito, David Keaton and Larry Wagoner
Original file name	n0259-mod-0907.pdf
Notes	Revision of N0259

Revised draft language-specific annex for the programming language C

Date	07-Sept-2010
Contributed by	Larry Wagoner
Original file name	c_language_annex_052710.docx
Notes	Replaces N0259 Modified – Benito July 2010 Modified – Keaton and Benito Sept 2010

Language Specific Vulnerability Outline

C. Skeleton template for use in proposing language specific information for vulnerabilities

Every vulnerability description of Clause 6 of the main document should be addressed in the annex in the same order even if there is simply a notation that it is not relevant to the language in question.

C.1 Standards and terminology

C.1.1 Identification of standards and associated documents

ISO/IEC 9899:1999 — *Programming Languages—C*

ISO/IEC TR 24731-1:2007 — *Extensions to the C library — Part 1: Bounds-checking interfaces*

ISO/IEC TR 24731-2:2010 — *Extensions to the C library — Part 2: Dynamic Allocation Functions*

ISO/IEC 9899:1999/Cor. 1:2001 — *Programming languages —C*

ISO/IEC 9899:1999/Cor. 2:2004 — *Programming languages —C*

ISO/IEC 9899:1999/Cor. 3:2007 — *Programming languages —C*

Seacord, Robert C. *The CERT C Secure Coding Standard*. Boston: Addison-Wesley, 2008.

GNU Project. GCC Bugs “Non-bugs” http://gcc.gnu.org/bugs.html#nonbugs_c (2009).

C.1.2 General terminology

access: An execution-time action, to read or modify the value of an object. Where only one of two actions is meant, *read* or *modify*. Modify includes the case where the new value being stored is the same as the previous value. Expressions that are not evaluated do not access objects.

alignment: The requirement that objects of a particular type be located on storage boundaries with addresses that are particular multiples of a byte address.

argument:

actual argument: The expression in the comma-separated list bounded by the parentheses in a function call expression, or a sequence of preprocessing tokens in the comma-separated list bounded by the parentheses in a function-like macro invocation.

behaviour: An external appearance or action.

implementation-defined behaviour: The *unspecified behaviour* where each implementation documents how the choice is made. An example of implementation-defined behaviour is the propagation of the high-order bit when a signed integer is shifted right.

locale-specific behaviour: The behaviour that depends on local conventions of nationality, culture, and language that each implementation documents. An example, locale-specific behaviour is whether the `islower()` function returns true for characters other than the 26 lower case Latin letters.

undefined behaviour: The use of a non-portable or erroneous program construct or of erroneous data, for which the C standard imposes no requirements. Undefined behaviour ranges from ignoring the situation completely with unpredictable results, to behaving during translation or program execution in a documented manner characteristic of the environment (with or without the issuance of a diagnostic message), to terminating a translation or execution (with the issuance of a diagnostic message). An example of, undefined behaviour is the behaviour on integer overflow.

unspecified behaviour: The use of an unspecified value, or other behaviour where the C Standard provides two or more possibilities and imposes no further requirements on which is chosen in any instance. For example, unspecified behaviour is the order in which the arguments to a function are evaluated.

bit: The unit of data storage in the execution environment large enough to hold an object that may have one of two values. It need not be possible to express the address of each individual bit of an object.

byte: The addressable unit of data storage large enough to hold any member of the basic character set of the execution environment. It is possible to express the address of each individual byte of an object uniquely. A byte is composed of a contiguous sequence of bits, the number of which is implementation-defined. The least significant bit is called the low-order bit; the most significant bit is called the high-order bit.

character: An abstract member of a set of elements used for the organization, control, or representation of data.

single-byte character: The bit representation that fits in a byte.

multibyte character: The sequence of one or more bytes representing a member of the extended character set of either the source or the execution environment. The extended character set is a superset of the basic character set.

wide character: The bit representation that will fit in an object capable of representing any character in the current locale. The C Standard uses the type name `wchar_t` for this object.

correctly rounded result: The representation in the result format that is nearest in value, subject to the current rounding mode, to what the result would be given unlimited range and precision.

diagnostic message: The message belonging to an implementation-defined subset of the implementation's message output. The C Standard requires diagnostic messages for all constraint violations.

implementation: A particular set of software, running in a particular translation environment under particular control options, that performs translation of programs for, and supports execution of functions in, a particular execution environment.

implementation limit: The restriction imposed upon programs by the implementation.

memory location: Either an object of scalar¹ type, or a maximal sequence of adjacent bit-fields all having nonzero width. A bit-field and an adjacent non-bit-field member are in separate memory locations. The same applies to two bit-fields-fi, if one is declared inside a nested structure declaration and the other is not, or if the two

¹ Integer types, Floating types and Pointer types are collectively called *scalar* types in the C Standard.

are separated by a zero-length bit-field declaration, or if they are separated by a non-bit-field member declaration. It is not safe to concurrently update two bit-field-fi in the same structure if all members declared between them are also bit-fields, no matter what the sizes of those intervening bit-fields happen to be. For example a structure declared as

```
struct {
  char a;
  int b:5, c:11, :0, d:8;
  struct { int ee:8; } e;
}
```

contains four separate memory locations: The member *a*, and bit-fields *d* and *e*. *ee* are separate memory locations, and can be modified concurrently without interfering with each other. The bit-fields *b* and *c* together constitute the fourth memory location. The bit-fields *b* and *c* can't be concurrently modified, but *b* and *a*, can be concurrently modified.

object: The region of data storage in the execution environment, the contents of which can represent values. When referenced, an object may be interpreted as having a particular type.

parameter:

formal parameter: The object declared as part of a function declaration or definition that acquires a value on entry to the function, or an identifier from the comma-separated list bounded by the parentheses immediately following the macro name in a function-like macro definition.

recommended practice: A specification that is strongly recommended as being in keeping with the intent of the C Standard, but that may be impractical for some implementations.

runtime-constraint: A requirement on a program when calling a library function.

value: The precise meaning of the contents of an object when interpreted as having a specific type.

implementation-defined value: An unspecified value where each implementation documents how the choice for the value is selected.

indeterminate value: Is either an unspecified value or a trap representation.

unspecified value: The valid value of the relevant type where the C Standard imposes no requirements on which value is chosen in any instance. An unspecified value cannot be a trap representation.

trap representation: An object representation that need not represent a value of the object type.

C.1.2 General concepts

C.2 Obscure Language Features [BRS]

C.2.1 Terminology and features

C.2.2 Description of vulnerability

C is a relatively small language with a limited syntax set lacking many of the complex features of some other languages. Many of the complex features in C are not implemented as part of the language syntax, but rather implemented as library routines. As such, most of the available features in C are used relatively frequently.

Common use across a variety of languages may make some features less obscure. Because of the unstructured code that is frequently the result of using `goto`'s, the `goto` statement is frequently restricted, or even outright banned, in some C development environments. Even though the `goto` is encountered infrequently and the use of it considered obscure, because it is fairly obvious as to its purpose and since its use is common to many other languages, the functionality of it is easily understood by even the most junior of programmers.

The use of a combination of features adds yet another dimension. Particular combinations of features in C may be used rarely together or fraught with issues if not used correctly in combination. This can cause unexpected results and potential vulnerabilities.

C.2.3 Avoiding the vulnerability or mitigating its effects

- Organizations should specify coding standards that restrict or ban the use of features or combinations of features that have been observed to lead to vulnerabilities in the operational environment for which the software is intended.

C.2.4 Implications for standardization

Future standardization efforts should consider:
None

C.3 Unspecified Behaviour [BQF]

C.3.1 Terminology and features

Unspecified behaviour occurs where the C standard provides two or more possibilities but does not dictate which one is chosen. Unspecified behaviour also occurs when an unspecified value is used.

An *unspecified value* is a value that is valid for its type and where the C standard does not impose a choice on the value chosen. Many aspects of the C language result in unspecified behaviour.

C.3.2 Description of vulnerability

The C standard has documented, in Annex J.1, 54 instances of unspecified behaviour. Examples of unspecified behaviour are:

- The order in which the operands of an assignment operator are evaluated
- The order in which any side effects occur among the initialization list expressions in an initializer
- The layout of storage for function parameters

Reliance on a particular behaviour that is unspecified leads to portability problems because the expected behaviour may be different for any given instance. Many cases of unspecified behaviour have to do with the order of evaluation of subexpressions and side effects. For example, in the function call

```
f1(f2(x), f3(x));
```

the functions `f2` and `f3` may be called in any order possibly yielding different results depending on the order in which the functions are called.

C.3.3 Avoiding the vulnerability or mitigating its effects

- Do not rely on unspecified behaviour because the behaviour can change at each instance. Thus, any code that makes assumptions about the behaviour of something that is unspecified should be replaced to make it less reliant on a particular installation and more portable.

C.3.4 Implications for standardization

Future standardization efforts should consider:

None

C.4 Undefined Behaviour [EWF]

C.4.1 Terminology and features

Undefined behaviour is behaviour that results from using erroneous constructs and data.

C.4.2 Description of vulnerability

The C standard does not impose any requirements on undefined behaviour. Typical undefined behaviours include doing nothing, producing unexpected results, and terminating the program.

The C standard has documented, in Annex J.2, 191 instances of undefined behaviour that exist in C. One example of undefined behaviour occurs when the value of the second operand of the `/` or `%` operator is zero. This is generally not detectable through static analysis of the code, but could easily be prevented by a check for a zero divisor before the operation is performed. Leaving this behaviour as undefined lessens the burden on the implementation of the division and modulo operators.

Other examples of undefined behaviour are:

- Referring to an object outside of its lifetime
- The conversion to or from an integer type that produces a value outside of the range that can be represented
- The use of two identifiers that differ only in non-significant characters

Relying on undefined behaviour makes a program unstable and non-portable. While some cases of undefined behaviour may be consistent across multiple implementations, it is still dangerous to rely on them. Relying on undefined behaviour can result in errors that are difficult to locate and only present themselves under special circumstances. For example, accessing memory deallocated by `free()` or `realloc()` results in undefined behaviour, but it may work most of the time.

C.4.3 Avoiding the vulnerability or mitigating its effects

- Eliminate to the extent possible all cases of undefined behaviour from a program

C.4.4 Implications for standardization

Future standardization efforts should consider:

Making the declarations of undefined behaviour more definitive². The collection of undefined behaviour in Annex J.2 is well done with cross references to sections in the standard. Most of the entries are well defined, but the few

² This is actually being addressed in Annex L of the C Working Paper, see [WG 14/N1494](#).

C.19 Boundary Beginning Violation [XYX]**C.19.1 Terminology and features****C.19.2 Description of vulnerability**

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

In C, the subscript operator [] is defined such that $E1[E2]$ is identical to $(*((E1)+(E2)))$, so that in either representation, the value in location $(E1+E2)$ is returned. Because C does not perform bounds checking on arrays, the following code:

```
int foo(const int i) {
    int x[] = {0,0,0,0,0,0,0,0,0,0};
    return x[i];
}
```

would return whatever is in location $x[i]$ even if, say, i were equal to -5 (assuming that $x[-5]$ was still within the address space of the program). This could be sensitive information or even a return address, which if altered by changing the value of $x[-5]$, could change the program flow.

C.19.3 Avoiding the vulnerability or mitigating its effects

- Perform range checking before accessing an array since C does not perform bounds checking automatically. In the interest of speed and efficiency, range checking only needs to be done when it cannot be statically shown that an access outside of the array cannot occur.
- Use safer and more secure functions for string handling from the ISO TR24731-1, Extensions to the C library— Part 1: Bounds-checking interfaces. These are alternative string handling library functions to the existing Standard C Library. The functions verify that receiving buffers are large enough for the resulting strings being placed in them and ensure that resulting strings are null terminated. One implementation of these functions has been released as the Safe C Library.

C.19.4 Implications for standardization

Future standardization efforts should consider:

- Defining an array type that does automatic bounds checking.

C.20 Unchecked Array Indexing [XYZ]**C.20.1 Terminology and features****C.20.2 Description of vulnerability**

C does not perform bounds checking on arrays, so though arrays may be accessed outside of their bounds, the value returned is undefined and in some cases may result in a program termination. For example, in C the following code is valid, though, for example, if i has the value 10, the result is undefined:

```
int foo(const int i) {
```

```
int t;  
int x[] = {0,0,0,0,0};  
t = x[i];  
return t;  
}
```

The variable `t` will likely be assigned whatever is in the location pointed to by `x[10]` (assuming that `x[10]` is still within the address space of the program).

C.20.3 Avoiding the vulnerability or mitigating its effects

- Perform range checking before accessing an array since C does not perform bounds checking automatically. In the interest of speed and efficiency, range checking only needs to be done when it cannot be statically shown that an access outside of the array cannot occur.
- Use safer and more secure functions for string handling from the ISO TR24731-1, Extensions to the C library— Part 1: Bounds-checking interfaces. These are alternative string handling library functions to the existing Standard C Library. The functions verify that receiving buffers are large enough for the resulting strings being placed in them and ensure that resulting strings are null terminated. One implementation of these functions has been released as the Safe C Library.

C.20.4 Implications for standardization

Future standardization efforts should consider:

- Defining an array type that does automatic bounds checking.

C.21 Unchecked Array Copying [XYW]

C.21.1 Terminology and features

C.21.2 Description of vulnerability

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

In the interest of ease and efficiency, C library functions such as `memcpy(void * restrict s1, const void * restrict s2, size_t n)` and `memmove(void *s1, const void *s2, size_t n)` are used to copy the contents from one area to another. `memcpy()` and `memmove()` simply copy memory and no checks are made as to whether the destination area is large enough to accommodate the `n` units of data being copied. It is assumed that the calling routine has ensured that adequate space has been provided in the destination. Problems can arise when the destination buffer is too small to receive the amount of data being copied or if the indices being used for either the source or destination are not the intended indices.

C.21.3 Avoiding the vulnerability or mitigating its effects

- Perform range checking before calling a memory copying function such as `memcpy()` and `memmove()`. These functions do not perform bounds checking automatically. In the interest of speed and efficiency, range checking only needs to be done when it cannot be statically shown that an access outside of the array cannot occur.

C.21.4 Implications for standardization

Future standardization efforts should consider:

- Defining functions that contain an extra parameter in `memcpy()` and `memmove()` for the maximum number of bytes to copy. In the past, some have suggested that the size of the destination buffer be used as an additional parameter. Some critics state that this solution is easy to circumvent by simply repeating the parameter that was used for the number of bytes to copy as the parameter for the size of the destination buffer. This analysis and criticism is correct. What is needed is a failsafe check as to the maximum number of bytes to copy. There are several reasons for creating new functions with an additional parameter. This would make it easier for static analysis to eliminate those cases where the memory copy could not be a problem (such as when the maximum number of bytes is demonstrably less than the capacity of the receiving buffer). Manual analysis or more involved static analysis could then be used for the remaining situations where the size of the destination buffer may not be sufficient for the maximum number of bytes to copy. This extra parameter may also help in determining which copies could take place among objects that overlap. Such copying is undefined according to the C standard. It is suggested that safer versions of functions that include a restriction `max_n` on the number of bytes `n` to copy (e.g. `void *memcpy(void * restrict s1, const void * restrict s2, size_t n), const size_t max_n`) be added to the standard in addition to retaining the current corresponding functions (e.g. `memcpy(void * restrict s1, const void * restrict s2, size_t n)`). The additional parameter would be consistent with the copying function pairs that have already been created such as `strcpy()/strncpy()` and `strcat()/strncat()`. This would allow a safer version of memory copying functions for those applications that want to use them in to facilitate both safer and more secure code and more efficient and accurate static code reviews.⁴

C.22 Buffer Overflow [XZB]

C.22.1 Terminology and features

C.22.2 Description of vulnerability

C is a flexible and efficient language due to its rather lax restrictions on memory manipulations. Writing outside of a buffer can occur easily in C due to a miscalculation of the size of the buffer, a mistake in a loop termination condition or any of dozens of other ways. Egregious violations of a buffer size are often found during testing as crashes of the program occur. However, more subtle or input dependent overflows may go undetected in testing and later be exploited by attackers.

As with other languages, it is easy to overflow a buffer in C. The main difference is that C does not prevent or detect the occurrence automatically as is done in many other languages. For instance, consider:

```
int foo(const int n) {
    char buf[10];
    for (i=1; i++; i<=n)
        buf[i] = i + 0x40;
    return buf[n];
}
```

A value of 10 for `n` will write 0x50 to `buf[10]` which is one beyond the end of the array `buf` which starts at `buf[0]` and ends at `buf[9]`. Overflows where the amount of the overflow and the content can be manipulated by an attacker can cause the program to crash or execute logic that gives the attacker host access. For instance, the `gets()` function has been deprecated since there isn't a way stop a user from typing in a longer string than expected and overrunning a buffer. Consider:

```
int main()
```

⁴ This has been addressed by WG 14 in an optionally normative annex in the current working paper

```
{
  char buf[500];
  printf "Type something.\n");
  gets(buf);
  printf "You typed: %s\n", buf);

  return 0;
}
```

Typing in a string longer than 499 characters (1 less than the buffer length to account for the string null termination character) will cause the buffer to overflow. A well crafted string used as input to this program can cause execution of an attacker's malicious code.

C.22.3 Avoiding the vulnerability or mitigating its effects

- Validate all input values.
- Check any array index before use if there is a possibility the value could be outside the bounds of the array.
- Use length restrictive functions such as `strncpy()` instead of `strcpy()`.
- Use stack guarding add-ons to prevent overflows of stack buffers.
- Do not use the deprecated functions or other language features such as `gets()`.
- Be aware that the use of all of these preventive measures may still not be able to stop all buffer overflows from happening. However, the use of them can make it much rarer for a buffer overflow to occur and much harder to exploit it.
- Use alternative functions as specified in ISO/IEC TR 24731-1:2007 or TR 24731-2:2010. These Technical Reports provides alternative functions for the C Library (as defined in ISO/IEC 9899:1999) that promotes safer, more secure programming. The functions verify that output buffers are large enough for the intended result and return a failure indicator if they are not. Optionally, failing functions call a "runtime-constraint handle" to report the error. Data is never written past the end of an array. All string results are null terminated. In addition, the functions in ISO/IEC TR 24731-1:2007 are re-entrant: they never return pointers to static objects owned by the function. ISO/IEC TR 24731-1:2007 also contains functions that address insecurities with the C input-output facilities.

C.22.4 Implications for standardization

Future standardization efforts should consider:

- Deprecating less safe functions such as `strcpy()` and `strcat()` where a more secure alternative is available (an annex in the WG 14 working paper addresses this issue).
- Defining safer and more secure replacement functions such as `memncpy()` and `memncat()` to complement the `memcpy()` and `memcat()` functions (an annex in the WG 14 working paper addresses this issue).
- Adopting one of the Technical Reports on safer C library functions, Extensions to the C Library (TR 24731-1: Part I: *Bounds-checking interfaces* or TR 24731-2: Part II: *Dynamic allocation functions*, that have been developed by WG14.

C.23 Pointer Casting and Pointer Type Changes [HFC]

C.23.1 Terminology and features

C.23.2 Description of vulnerability

C allows casting the value of a pointer to and from another data type. These conversions can cause unexpected changes to pointer values.

Pointers in C refer to a specific type, such as integer. If `sizeof(int)` is 4 bytes, and `ptr` is a pointer to integers that contains the value `0x5000`, then `ptr++` would make `ptr` equal to `0x5004`. However, if `ptr` were a pointer to `char`, then `ptr++` would make `ptr` equal to `0x5001`. It is the difference due to data sizes coupled with conversions between pointer data types that cause unexpected results and potential vulnerabilities. Due to arithmetic operations, pointers may not maintain correct memory alignment or may operate upon the wrong memory addresses.

C.23.3 Avoiding the vulnerability or mitigating its effects

- Maintain the same type to avoid errors introduced through conversions.
- Heed compiler warnings that are issued for pointer conversion instances. The decision may be made to avoid all conversions so any warnings must be addressed. Note that casting into and out of "void*" pointers will most likely not generate a compiler warning as this is valid in C.

C.23.4 Implications for standardization

Future standardization efforts should consider:
None

C.24 Pointer Arithmetic [RVG]

C.24.1 Terminology and features

C.24.2 Description of vulnerability

When performing pointer arithmetic in C, the size of the value to add to a pointer is automatically scaled to the size of the type of the pointed-to object. For instance, when adding a value to the byte address of a 4-byte integer, the value is scaled by a factor 4 and then added to the pointer. The effect of this scaling is that if a pointer `P` points to the `i`-th element of an array object, then `(P) + N` will point to the `i+n`-th element of the array. Failing to understand how pointer arithmetic works can lead to miscalculations that result in serious errors, such as buffer overflows.

The following example will illustrate arithmetic in C involving a pointer and how the operation is done relative to the size of the pointer's target. Consider the following code snippet:

```
int buf[5];
int *buf_ptr = buf;
```

where the address of `buf` is `0x1234`. Adding 1 to `buf_ptr` will result in `buf_ptr` being equal to `0x1238` on a host where an `int` is 4 bytes. `buf_ptr` will then contain the address of `buf[1]`. Not realizing that address operations will be in terms of the size of the object being pointed to can lead to address miscalculations and undefined behaviour.

C.24.3 Avoiding the vulnerability or mitigating its effects

- Consider an outright ban on pointer arithmetic due to the error prone nature of pointer arithmetic.
- Avoid the common pitfalls of pointer arithmetic. For instance, in checking the end of an array, the following method can be used:


```

int buf[INTBUFSIZE];
int *buf_ptr = buf;

while (havedata() && (buf_ptr < &buf[INTBUFSIZE])) /* buf[INTBUFSIZE]
                                                    is the address of the element
                                                    following the buf array */
{
    *buf_ptr++ = parseint(getdata());
}

```

C.24.4 Implications for standardization in

Future standardization efforts should consider:

- Restrictions on pointer arithmetic that could eliminate common pitfalls. Pointer arithmetic is error prone and the flexibility that it offers is useful, but some of the flexibility is simply a shortcut that if restricted could lessen the chance of a pointer arithmetic based error.

C.25 Null Pointer Dereference [XYH]

C.25.1 Terminology and features

C.25.2 Description of vulnerability

C allows memory to be dynamically allocated primarily through the use of `malloc()`, `calloc()`, and `realloc()`. Each will return the address to the allocated memory. Due to a variety of situations, the memory allocation may not occur as expected and a null pointer will be returned. Other operations or faults in logic can result in a memory pointer being set to null. Using the null pointer as though it pointed to a valid memory location can cause a segmentation fault and other unanticipated situations.

Space for 10000 integers can be dynamically allocated in C in the following way:

```

int *ptr = malloc(10000*sizeof(int)); /* allocate space for 10000 ints
*/

```

`malloc()` will return the address of the memory allocation or a null pointer if insufficient memory is available for the allocation. It is good practice after the attempted allocation to check whether the memory has been allocated via an `if` test against `NULL`:

```

if (ptr != NULL) /* check to see that the memory could be allocated */

```

Memory allocations usually succeed, so neglecting this test and using the memory will usually work. That is why neglecting the null test will frequently go unnoticed. An attacker can intentionally create a situation where the memory allocation will fail leading to a segmentation fault.

Faults in logic can cause a code path that will use a memory pointer that was not dynamically allocated or after memory has been deallocated and the pointer was set to null as good practice would indicate.

C.25.3 Avoiding the vulnerability or mitigating its effects

- Check whether a pointer is null before dereferencing it. As this can be overly extreme in many cases (such as in a `for` loop that performs operations on each element of a large segment of memory), judicious

checking of the value of the pointer at key strategic points in the code is recommended.

C.25.4 Implications for standardization

Future standardization efforts should consider:

None

C.26 Dangling Reference to Heap [XYK]

C.26.1 Terminology and features

C.26.2 Description of vulnerability

C allows memory to be dynamically allocated primarily through the use of `malloc()`, `calloc()`, and `realloc()`. C allows a considerable amount of freedom in accessing the dynamic memory. Pointers to the dynamic memory can be created to perform operations on the memory. Once the memory is no longer needed, it can be released through the use of `free()`. However, freeing the memory does not prevent the use of the pointers to the memory and issues can arise if operations are performed after memory has been freed.

Consider the following segment of code:

```
int foo() {
    int *ptr = malloc (100*sizeof(int));/* allocate space for 100 integers*/
    if (ptr != NULL) /* check to see that the memory could be allocated */
    {
        ... /* perform some operations on the dynamic memory */
        free (ptr); /* memory is no longer needed, so free it */
        ... /* program continues performing other operations */
        ptr[0] = 10;/* ERROR - memory is being used after it has been
released */
        ...
    }
    ...
}
```

The use of memory in C after it has been freed is undefined. Depending on the execution path taken in the program, freed memory may still be free or may have been allocated via another `malloc()` or other dynamic memory allocation. If the memory that is used is still free, use of the memory may be unnoticed. However, if the memory has been reallocated, altering of the data contained in the memory can result in data corruption. Determining that a dangling memory reference is the cause of a problem and locating it can be difficult.

Setting and using another pointer to the same section of dynamically allocated memory can also lead to undefined behaviour. Consider the following section of code:

```
int foo() {
    int *ptr = malloc (100*sizeof(int));/* allocate space for 100 integers
*/
    if (ptr != NULL) /* check to see that the memory could be allocated */
    {
        int ptr2 = &ptr[10]; /* set ptr2 to point to the 10th element of the
allocated memory */
        ... /* perform some operations on the dynamic memory */
        free (ptr); /* memory is no longer needed, so free it */
    }
}
```

```

    ptr = NULL; /* set ptr to NULL to prevent ptr from being used again */
    ...      /* program continues performing other operations */
    ptr2[0] = 10; /* ERROR - memory is being used after it has been released
via ptr2 */
    ...
}
return (0);
}

```

Dynamic memory was allocated via a `malloc()` and then later in the code, `ptr2` was used to point to an address in the dynamically allocated memory. After the memory was freed using `free(ptr)` and the good practice of setting `ptr` to `NULL` was followed to avoid a dangling reference by `ptr` later in the code, a dangling reference still existed using `ptr2`.

C.26.3 Avoiding the vulnerability or mitigating its effects

- Set a freed pointer to null immediately after a `free()` call, as illustrated in the following code:

```

free (ptr);
ptr = NULL;

```
- Do not create and use additional pointers to dynamically allocated memory.
- Only reference dynamically allocated memory using the pointer that was used to allocate the memory.

C.26.4 Implications for standardization

Future standardization efforts should consider:

- Modifying the library `free(void *ptr)` so that it sets `ptr` to `NULL` to prevent reuse of `ptr`.

C.27 Templates and Generics [SYM]

Does not apply to C.

C.28 Inheritance [RIP]

Does not apply to C.

C.29 Initialization of Variables [LAV]

C.29.0 Status and history

C.29.1 Terminology and features

C.29.2 Description of vulnerability

Local, automatic variables can assume unexpected values if they are used before they are initialized. The C Standard specifies, "If an object that has automatic storage duration is not initialized explicitly, its value is indeterminate". In the common case, on architectures that make use of a program stack, this value defaults to whichever values are currently stored in stack memory. While uninitialized memory often contains zeros, this is not guaranteed. Consequently, uninitialized memory can cause a program to behave in an unpredictable or unplanned manner and may provide an avenue for attack.

Assuming that an uninitialized variable is 0 can lead to unpredictable program behaviour when the variable is initialized to a value other than 0.

Many implementations will issue a diagnostic message indicating that a variable was not initialized.

C.29.3 Avoiding the vulnerability or mitigating its effects

- Heed compiler warning messages about uninitialized variables. These warnings should be resolved as recommended to achieve a clean compile at high warning levels.
- Do not use memory allocated by functions such as `malloc()` before the memory is initialized as the memory contents are indeterminate.

C.29.4 Implications for standardization

Future standardization efforts should consider:
None

C.30 Wrap-around Error [XYY]

C.30.1 Terminology and features

C.30.2 Description of vulnerability

Given the limited size of any computer data type, continuously adding one to the data type eventually will cause the value to go from a the maximum possible value to a small value. C permits this to happen without any detection or notification mechanism.

C is often used for bit manipulation. Part of this is due to the capabilities in C to mask bits and shift them. Another part is due to the relative closeness C has to assembly instructions. Manipulating bits on a signed value can inadvertently change the sign bit resulting in a number potentially going from a large positive value to a large negative value.

For example, consider the following code for a `short int` containing 16 bits:

```
int foo( short int i ) {  
    i++;  
    return i;  
}
```

Calling `foo` with the value of 32767 would cause undefined behaviour, such as wrapping to -32768. Manipulating a value in this way can result in unexpected results such as overflowing a buffer.

In C, bit shifting by a value that is greater than the size of the data type or by a negative number is undefined. The following code, where a `int` is 16 bits, would be undefined when `j` is greater than or equal to 16 or negative:

```
int foo( int i, const int j ) {  
    return i>>j;  
}
```

C.30.3 Avoiding the vulnerability or mitigating its effects

- Be aware that any of the following operators have the potential to wrap in C:

a + b	a - b	a * b	a++	a--	a += b
a -= b	a *= b	a << b	a >> b	-a	

- Use defensive programming techniques to check whether an operation will overflow or underflow the receiving data type. These techniques can be omitted if it can be shown at compile time that overflow or underflow is not possible.
- Only conduct bit manipulations on unsigned data types. The number of bits to be shifted by a shift operator should lie between 1 and (n-1), where n is the size of the data type.

C.30.4 Implications for standardization

Future standardization efforts should consider:

None

C.31 Sign Extension Error [XZI]

Does not apply to C, since instead of conversion routines, C uses direct casts and implicit conversions. This allows the compiler to pick the correct signedness.

C.32 Operator Precedence/Order of Evaluation [JCW]

C.32.1 Terminology and features

C.32.2 Description of vulnerability

The order in which an expression is evaluated can drastically alter the result of the expression. The order of evaluation of the operands in C is clearly defined, but misinterpretations by programmers can lead to unexpected results.

Consider the following:

```
int foo(short int a, short int b) {
    if (a | 0x7 == b)
        ...
}
```

designed to mask off and test the lower three bits of “a” for equality to “b”. However, due to the precedence rules in C, the effect of this expression is to perform the “0x7 == b” and then bitwise OR that with “a” which may or may not be the expected answer.

C.32.3 Avoiding the vulnerability or mitigating its effects

- Use parentheses generously to avoid any uncertainty or lack of portability in the order of evaluation of an expression. If parenthesis were used in the previous example, as in:

```
int foo(short int a, short int b) {
    if ((a | 0x7) == b)
        ...
}
```

the order of the evaluation would be clear.

C.32.4 Implications for standardization

Future standardization efforts should consider:

- Creating a few standardized precedence orders. Standardizing on a few precedence orders will help to eliminate the confusing intricacies that exist between languages. This would not affect current languages as altering precedence orders in existing languages is too onerous. However, this would set a basis for the future as new languages are created and adopted. Stating that a language uses "ISO precedence order A" would be useful rather than having to spell out the entire precedence order that differs in a conceptually minor way from some other languages, but in a major way when programmers attempt to switch between languages.

C.33 Side-effects and Order of Evaluation [SAM]

C.33.1 Terminology and features

C.33.2 Description of vulnerability

C allows expressions to have side effects. If two or more side effects modify the same expression as in:

```
int v[10];
int i;
/* ... */
i = v[i++];
```

the behaviour is undefined and this can lead to unexpected results. Either the "i++" is performed first or the assignment "i=v[i]" is performed first. Because the order of evaluation can have drastic effects on the functionality of the code, this can greatly impact portability.

There are several situations in C where the order of evaluation of subexpressions or the order in which side effects take place is unspecified including:

- The order in which the arguments to a function are evaluated (C99, Section 6.5.2.2, "Function calls").
- The order of evaluation of the operands in an assignment statement (C99, Section 6.5.16, "Assignment operators").
- The order in which any side effects occur among the initialization list expressions is unspecified. In particular, the evaluation order need not be the same as the order of subobject initialization (C99, Section 6.7.8, "Initialization").

Because these are unspecified behaviours, testing may give the false impression that the code is working and portable, when it could just be that the values provided cause evaluations to be performed in a particular order that causes side effects to occur as expected.

C.33.3 Avoiding the vulnerability or mitigating its effects

- Expressions should be written so that the same effects will occur under any order of evaluation that the C standard permits since side effects can be dependent on an implementation specific order of evaluation.

C.33.4 Implications for standardization

Future standardization efforts should consider:

None

C.34 Likely Incorrect Expression [KOA]

C.34.1 Terminology and features

C.34.2 Description of vulnerability

C has several instances of operators which are similar in structure, but vastly different in meaning. This is so common that the C example of confusing the Boolean operator “==” with the assignment “=” is frequently cited as an example among programming languages. Using an expression that is technically correct, but which may just be a null statement can lead to unexpected results.

C also provides a lot of freedom in constructing statements. This freedom, if misused, can result in unexpected results and potential vulnerabilities.

The flexibility of C can obscure the intent of a programmer. Consider:

```
int x,y;
/* ... */
if (x = y)
{
    /* ... */
}
```

A fair amount of analysis may need to be done to determine whether the programmer intended to do an assignment as part of the `if` statement (perfectly valid in C) or whether the programmer made the common mistake of using an “=” instead of a “==”. In order to prevent this confusion, it is suggested that any assignments in contexts that are easily misunderstood be moved outside of the Boolean expression. This would change the example code to:

```
int x,y;
/* ... */
x = y;
if (x == 0)
{
    /* ... */
}
```

This would clearly state what the programmer meant and that the assignment of `y` to `x` was intended.

Programmers can easily get in the habit of inserting the “;” statement terminator at the end of statements. However, inadvertently doing this can drastically alter the meaning of code, even though the code is valid as in the following example:

```
int a,b;
/* ... */
if (a == b); /* the semi-colon will make this a null statement */
{
    /* ... */
}
```

Because of the misplaced semi-colon, the code block following the `if` will always be executed. In this case, it is extremely likely that the programmer did not intend to put the semi-colon there.

C.34.3 Avoiding the vulnerability or mitigating its effects

- Simplify statements with interspersed comments to aid in accurately programming functionality and help future maintainers understand the intent and nuances of the code. The flexibility of C permits a programmer to create extremely complex expressions.
- Assignments embedded within other statements can be potentially problematic. Each of the following would be clearer and have less potential for problems if the embedded assignments were conducted outside of the expressions:

```
int a,b,c,d;
/* ... */
if ((a == b) || (c = (d-1))) /* the assignment to c may not occur */
/* if a is equal to b */
```

or:

```
int a,b,c;
/* ... */
foo (a=b, c);
```

Each is a valid C statement, but each may have unintended results.

- Null statements should have a source line of their own. This, combined with enforcement by static analysis, would make clearer the intention that the statement was meant to be a null statement.

C.34.4 Implications for standardization

Future standardization efforts should consider:

None

C.35 Dead and Deactivated Code [XYQ]

C.35.1 Terminology and features

C.35.2 Description of vulnerability

As with any programming language that contains branching statements, C programs can potentially contain dead code. It is of concern primarily since dead code may reveal a logic flaw or an unintentional mistake on the part of the programmer. Sometimes statements can be inserted in C programs as defensive programming such as adding a default case to a switch statement even though the expectation is that the default can never be reached – until through some twist of logic or through modifications to the code the notifying error message reveals the surprising event. These types of defensive statements may be able to be shown to be computationally impossible and thus are dead code. Those are not the focus. The focus is on those statements which are not defensive and which are unreachable. It is impossible to identify all such cases and therefore only those which are blatant and that indicate deeper issues of flawed logic may be able to be identified and removed.

C uses some operators that can be confused with other operators. For instance, the common mistake of using an assignment operator in a Boolean test as in:

```
int a,b;
/* ... */
```



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

can cause portions of code to become dead code since unless `b` can contain the value 0, the `else` portion of the `if` statement cannot be reached.

C.35.3 Avoiding the vulnerability or mitigating its effects

- Eliminate dead code to the extent possible from C programs.
- Use compilers and analysis tools to assist in identifying unreachable code.
- Use `/**/` comment syntax instead of `/*...*/` comment syntax to avoid the inadvertent commenting out sections of code.
- Delete deactivated code from programs due to the possibility of accidentally activating it.

C.35.4 Implications for standardization

Future standardization efforts should consider:
None

C.36 Switch Statements and Static Analysis [CLL]

C.36.1 Terminology and features

C.36.2 Description of vulnerability

Because of the way in which the switch-case statement in C is structured, it can be relatively easy to unintentionally omit the `break` statement between cases causing unintended execution of statements for some cases.

C contains a `switch` statement of the form:

```
char abc;
/* ... */
switch (abc)
{
  case 1:
    sval = "a";
    break;
  case 2:
    sval = "b";
    break;
  case 3:
    sval = "c";
    break;
  default:
    printf ("Invalid selection\n");
}
```

If there isn't a default case and the switched expression doesn't match any of the cases, then control simply shifts to the next statement after the switch statement block. Unintentionally omitting a `break` statement between two cases will cause subsequent cases to be executed until a `break` or the end of the switch block is reached. This could cause unexpected results.

C.36.3 Avoiding the vulnerability or mitigating its effects

- Only a direct fall through should be allowed from one case to another. That is, every nonempty case statement should be terminated with a `break` statement as illustrated in the following example:

```
int i;
/* ... */
switch (i)
{
  case 1:
  case 2:
    i++;      /* fall through from case 1 to 2 is permitted */
    break;
  case 3:
    j++;
  case 4:
    /* fall through from case 3 to 4 is not permitted */
    /* as it is not a direct fall through due to the */
    /* j++ statement */
}
```

- All `switch` statements should have a default value if only to indicate that there could exist a case that was unanticipated and thought impossible by the developers. The only exception is for switches on an enumerated type where all possible values can be exhausted. Even in the case of enumerated types, it is suggested that a default be inserted in anticipation of possible code changes to the enumerated type.

C.36.4 Implications for standardization

Future standardization efforts should consider:

- Defining a “fallthru” construct that will explicitly bind multiple switch cases together and eliminate the need for the `break` statement. The default would be for a case to break instead of falling through to the next case. Granted this is a major shift in concept, but if it could be accomplished, less unintentional errors would occur.

C.37 Demarcation of Control Flow [EOJ]

C.37.1 Terminology and features

A *block-structured language* is a language that has a syntax for enclosing structures between bracketed keywords, such as an `if` statement bracketed by `if` and `endif`, as in FORTRAN, or a code section bracketed by `BEGIN` and `END`, as in PL/1.

A *comb-structured language* is a language that has an ordered set of keywords to define separate sections within a block, analogous to the multiple teeth or prongs in a comb separating sections of the comb. For example, in Ada, a block is a 4-pronged comb with keywords `declare`, `begin`, `exception`, `end`, and the `if` statement in Ada is a 4-pronged comb with keywords `if`, `then`, `else`, `end if`.

C.37.2 Description of vulnerability

C is a block-structured language, while languages such as Ada and Pascal are comb-structured languages. Therefore, it may not be readily apparent which statements are part of a loop construct or an `if` statement.

Consider the following section of code:

```
int foo(int a, const int *b) {
```

```

int i=0;

/* ... */
a = 0;
for (i=0; i<10; i++);
{
    a = a + b[i];
}
}

```

At first it may appear that `a` will be a sum of the numbers `b[0]` to `b[9]`. However, even though the code is structured so that the “`a = a + b[i]`” code is structured to appear within the `for` loop, the “`;`” at the end of the `for` statement causes the loop to be on a null statement (the “`;`”) and the “`a = a + b[i];`” statement to only be executed once. In this case, this mistake may be readily apparent during development or testing. More subtle cases may not be as readily apparent leading to unexpected results.

If statements in C are also susceptible to control flow problems since there isn't a requirement in C for there to be an `else` statement for every `if` statement. An `else` statement in C always belong to the most recent `if` statement without an `else`. However, the situation could occur where it is not readily apparent to which `if` statement an `else` due to the way the code is indented or aligned.

C.37.3 Avoiding the vulnerability or mitigating its effects

- Enclose the bodies of `if`, `else`, `while`, `for`, etc. in braces. This will reduce confusion and potential problems when modifying the software. For example:

```

int a,b,i;

/* ... */

if (i = 10)
{
    a = 5;          /* this is correct */
    b = 10;
}
else
    a = 10;        /* this is incorrect -- the assignments to b */
                  /* were added later and were expected to */
    b = 5;         /* be part of the if and else and indented */
                  /* as such, but did not become part of the else */

```

- Use a final `else` statement or a comment stating why the final `else` isn't necessary in all `if` and `else if` statements.

C.37.4 Implications for standardization

Future standardization efforts should consider:

None

C.38 Loop Control Variables [TEX]

C.38.1 Terminology and features

C.38.2 Description of vulnerability

C allows the modification of loop control variables within a loop. Though this is usually not considered good programming practice as it can cause unexpected problems, the flexibility of C expects the programmer to use this capability responsibly.

Since the modification of a loop control variable within a loop is infrequently encountered, reviewers of C code may not expect it and hence miss noticing the modification. Modifying the loop control variable can cause unexpected results if not carefully done. In C, the following is valid:

```
int a,i;

for (i=1; i<10; i++)
{
    ...
    if (a > 7)
        i = 10;
    ...
}
```

which would cause the `for` loop to exit once `a` is greater than 7 regardless of the number of loops that have occurred.

C.38.3 Avoiding the vulnerability or mitigating its effects

- Do not modify a loop control variable within a loop. Even though the capability exists in C, it is still considered to be a poor programming practice.

C.38.4 Implications for standardization

Future standardization efforts should consider:

- Defining an identifier type for loop control that cannot be modified by anything other than the loop control construct would be a relatively minor addition to C that could make C code safer and encourage better structured programming.

C.39 Off-by-one Error [XZH]

C.39.1 Terminology and features

C.39.2 Description of vulnerability

Arrays are a common place for off by one errors to manifest. In C, arrays are indexed starting at 0, causing the common mistake of looping from 0 to the size of the array as in:

```
int foo() {
    int a[10];
    int i;
    for (i=0, i<=10, i++)
        ...
    return (0);
}
```

Strings in C are also another common source of errors in C due to the need to allocate space for and account for

the string sentinel value. A common mistake is to expect to store an n length string in an n length array instead of length $n+1$ to account for the sentinel `'\0'`. Interfacing with other languages that do not use sentinel values in strings can also lead to an off by one error.

C does not flag accesses outside of array bounds, so an off by one error may not be as detectable in C as in some other languages. Several good and freely available tools for C can be used to help detect accesses beyond the bounds of arrays that are caused by an off by one error. However, such tools will not help in the case where only a portion of the array is used and the access is still within the bounds of the array.

Looping one more or one less is usually detectable by good testing. Due to the structure of the C language, this may be the main way to avoid this vulnerability. Unfortunately some cases may still slip through the development and test phase and manifest themselves during operational use.

C.39.3 Avoiding the vulnerability or mitigating its effects

- Use careful programming, testing of border conditions and static analysis tools to detect off by one errors in C.

C.39.4 Implications for standardization

Future standardization efforts should consider:

None

C.40 Structured Programming [EWD]

C.40.1 Terminology and features

C.40.2 Description of vulnerability

It is as easy to write structured programs in C as it is not to. C contains the `goto` statement, which can create unstructured code. Also, C has `continue`, `break`, and `return` that can create a complicated control flow, when used in an undisciplined manner. Spaghetti code can be more difficult for C static analyzers to analyze and is sometimes used on purpose to intentionally obfuscate the functionality of software. Code that has been modified multiple times by an assortment of programmers to add or remove functionality or to fix problems can be prone to become unstructured.

Because unstructured code in C can cause problems for analyzers (both automated and human) of code, problems with the code may not be detected as readily or at all as would be the case if the software was written in a structured manner.

C.40.3 Avoiding the vulnerability or mitigating its effects

- Write clear and concise structured code to make code as understandable as possible.
- Restrict the use of `goto`, `continue`, `break` and `return` to encourage more structured programming.
- Encourage the use of a single exit point from a function. At times, this guidance can have the opposite effect, such as in the case of an `if` check of parameters at the start of a function that requires the remainder of the function to be encased in the `if` statement in order to reach the single exit point. If, for example, the use of multiple exit points can arguably make a piece of code clearer, then they should be used. However, the code should be able to withstand a critique that a restructuring of the code would have made the need for multiple exit points unnecessary.

C.40.4 Implications for standardization

Future standardization efforts should consider:

- Deprecating the `goto` statement. The use of the `goto` construct is often spotlighted as the antithesis of good structured programming. Though its deprecation will not instantly make all C code structured, deprecating the `goto` and leaving in place the restricted `goto` variations (e.g. `break` and `continue`) and possibly adding other restricted `goto`'s could assist in encouraging safer and more secure C programming in general.
-

C.41 Passing Parameters and Return Values [CSJ]

C.41.1 Terminology and features

C.41.2 Description of vulnerability

At times, it is useful to interface a C program with routines written in other languages. Other languages may have different data types, storage orders or parameter passing semantics. These differences in interfacing with other languages can lead to unexpected interpretations or manipulations of data.

C only passes parameters by value. That is, the receiving function will get the value of the parameter. Call by reference can be achieved by passing a reference as a value. Interfacing with another language, such as Fortran, that uses call by reference can yield some surprising results. Therefore, the addresses of the arguments must be passed when calling a Fortran subroutine from C. There are many other major and minor issues in interfacing to other languages all of which can lead to unexpected results and even potential vulnerabilities. For example, arrays in C are stored in row major order (last index varies fastest) whereas Fortran stores arrays in column major order (first index varies fastest). Other issues are minor annoyances, such as the inability of C to be able to pass a constant as a parameter to a Fortran subroutine since there isn't an address to pass (that is, `&7`) to satisfy the call by reference expectation.

C.41.3 Avoiding the vulnerability or mitigating its effects

- Use caution when interfacing with other languages as this can be error prone.
- Use interface packages that are available for many language combinations which can assist in avoiding some problems in interfacing. Even with an interface package, there will likely still be some issues that need to be addressed for a successful interface.
- Conduct additional rigorous testing on sections of code that interface with other languages.

C.41.4 Implications for standardization

Future standardization efforts should consider:

- Defining a standardized interface package for interfacing C with many of the top programming languages and a reciprocal package should be developed of the other top languages to interface with C.
-

C.42 Dangling References to Stack Frames [DCM]

C.42.1 Terminology and features

C.42.2 Description of vulnerability

C allows the address of a variable to be stored in a variable. Should this variable's address be, for example, the

address of a local variable that was part of a stack frame, then using the address after the local variable has been deallocated can yield unexpected behaviour as the memory will have been made available for further allocation and may indeed be allocated for some other use. Any use of perishable memory after it has been deallocated can lead to unexpected results.

C.42.3 Avoiding the vulnerability or mitigating its effects

- Do not assign the address of an object to any entity which persists after the object has ceased to exist. This is done in order to avoid the possibility of a dangling reference. Once the object ceases to exist, then so will the stored address of the object preventing accidental dangling references.
- Pointers should be assigned the null-pointer value before executing a return for any block-local addresses that have been stored in longer-lived storage.

C.42.4 Implications for standardization

Future standardization efforts should consider:
None

C.43 Subprogram Signature Mismatch [OTR]

C.43.1 Terminology and features

C.43.2 Description of vulnerability

Functions in C may be called with more or less than the number of parameters the receiving function expects. However, most C compilers will generate a warning or an error about this situation. If the number of arguments does not equal the number of parameters, the behaviour is undefined. This can lead to unexpected results when the count or types of the parameters differs from the calling to the receiving function. If too few arguments are sent to a function, then the function could still pop the expected number of arguments from the stack leading to unexpected results.

C allows a variable number of arguments in function calls. A good example of an implementation of this is the `printf()` function. This is specified in the function call by terminating the list of parameters with an ellipsis (`, ...`). After the comma, no information about the number or types of the parameters is supplied. This can be a useful feature for situations such as `printf()`, but the use of this feature outside of special situations can be the basis for vulnerabilities.

Functions may or may not be defined with a function definition. The function definition may or may not contain a parameter type list. If a function that accepts a variable number of arguments is defined without a parameter type list that ends with the ellipsis notation, the behaviour is undefined.

If the calling and receiving functions differ in the type of parameters, C will, if possible, do an implicit conversion such as the call to `sqrt()` that expects a double:

```
double sqrt(double)
```

the call:

```
root2 = sqrt(2);
```

coerces the integer 2 into the double value 2.0.

C.43.3 Avoiding the vulnerability or mitigating its effects

- Use a function prototype to declare a function with its expected parameters to allow the compiler to check for a matching count and types of the parameters. The prototype contains just the name of the function and its parameters without the body of code that would normally follow.
- Do not use the variable argument feature except in rare instances. The variable argument feature such as is used in `printf()` is difficult to use in a type safe manner.

C.43.4 Implications for standardization

Future standardization efforts should consider:

None

C.44 Recursion [GDL]**C.44.1 Terminology and features****C.44.2 Description of vulnerability**

C permits recursive calls both directly and indirectly through any chain of other functions. However, recursive functions must be implemented carefully in C, C does not have protective mechanisms that could avert serious problems such as an overly large consumption of resources or an overrun of buffers. Since C is frequently cited for its high performance efficiency, the use of recursion in C can be counter to this as recursion can be inefficient both in execution time and memory usage. Some of the modern compilers perform tail-call optimization to make recursion efficient and resource-friendly.

As with many languages, the high consumption of resources for recursive calls can apply to C. It is difficult to predict the complete range of values that a recursive function can execute that will lead to a manageable consumption of resources. Part of this difficulty is that the range of values can change depending on the current load of the host. Manipulation of the input values to a recursive function can result in an intentional exhaustion of system resources leading to a denial of service.

C.44.3 Avoiding the vulnerability or mitigating its effects

- Only use recursion in rare instances. Although recursion can shorten programs considerably, there is a high performance penalty which is contrary to the usual high efficiency of C.
- Only use recursion if it can be proven that adequate resources exist to support the maximum level of recursion possible.

C.44.4 Implications for standardization

Future standardization efforts should consider:

None

C.45 Returning Error Status [NZN]**C.45.1 Terminology and features****C.45.2 Description of vulnerability**

C provides the include file `<errno.h>` that defines the macros `EDOM`, `EILSEQ` and `ERANGE`, which expand to integer constant expressions with type `int`, distinct positive values and which are suitable for use in `#if` preprocessing directives. C also provides the integer `errno` that can be set to a nonzero value by any library function (if the use of `errno` is not documented in the description of the function in the C Standard, `errno` could be used whether or not there is an error). Though these values are defined, inconsistencies in responding to error conditions can lead to vulnerabilities.

C.45.3 Avoiding the vulnerability or mitigating its effects

- Check the returned error status upon return from a function. The C standard library functions provide an error status as the return value and sometimes in an additional global error value.
- Set `errno` to zero before a library function call in situations where a program intends to check `errno` before a subsequent library function call.
- Use `errno_t` to make it readily apparent that a function is returning an error code. Often a function that returns an `errno` error code is declared as returning a value of type `int`. Although syntactically correct, it is not apparent that the return code is an `errno` error code. TR 24731-1 introduced the new type `errno_t` in `<errno.h>` that is defined to be type `int`.

C.45.4 Implications for standardization

Future standardization efforts should consider:

- Joining with other languages in developing a standardized set of mechanisms for detecting and treating error conditions so that all languages to the extent possible could use them. Note that this does not mean that all languages should use the same mechanisms as there should be a variety (e.g. label parameters, auxiliary status variables), but each of the mechanisms should be standardized.

C.46 Termination Strategy [REU]

C.46.1 Terminology and features

C.46.2 Description of vulnerability

Choosing when and where to exit is a design issue, but choosing how to perform the exit may result in the host being left in an unexpected state. C provides several ways of terminating a program including `exit()`, `_Exit()`, and `abort()`. A return from the initial call to the `main` function is equivalent to calling the `exit()` function with the value returned by the `main` function as its argument (this is if the return type of the `main` function is a type compatible with `int`, otherwise the termination status returned to the host environment is unspecified) or simply reaching the `}"` that terminates the `main` function returns a value of 0.

All of the termination strategies in C have undefined, unspecified, and/or implementation defined behaviour associated with them. For example, if more than one call to the `exit()` function is executed by a program, the behaviour is undefined. The amount of clean-up that occurs upon termination such as the removal of temporary files or the flushing of buffers varies and may be implementation defined.

A call to `exit()` or `_Exit()` will terminate a program normally. Abnormal program termination will occur when `abort()` is used to exit a program (unless the signal `SIGABRT` is caught and the signal handler does not return). Unlike a call to `exit()`, when either `_Exit()` or `abort()` are used to terminate a program, it is implementation defined as to whether open streams with unwritten buffered data are flushed, open streams are closed, or temporary files are removed. This can leave a system in an unexpected state.

C provides the function `atexit()` that allows functions to be registered so that at normal program termination, the registered functions will be executed to perform desired functions. C99 requires the capability to register *at least* 32 functions. Implementations expecting more than 32 registered functions may yield unexpected results.

C.46.3 Avoiding the vulnerability or mitigating its effects

- Use a return from the `main()` program as it is the cleanest way to exit a C program.
- Use `exit()` to quickly exit from a deeply nested function.
- Use `abort()` in situations where an abrupt halt is needed. If `abort()` is necessary, the design should protect critical data from being exposed after an abrupt halt of the program.
- Become familiar with the undefined, unspecified and/or implementation aspects of each of the termination strategies.

C.46.4 Implications for standardization

Future standardization efforts should consider:

- Since fault handling and exiting of a program is common to all languages, it is suggested that common terminology such as the meaning of fail safe, fail hard, fail soft, etc. along with a core API set such as `exit`, `abort`, etc. be standardized and coordinated with other languages.

C.47 Extra Intrinsic [LRM]

Does not apply to C.

C.48 Type-breaking Reinterpretation of Data [AMV]

C.48.1 Terminology and features

C.48.2 Description of vulnerability

The primary way in C that a reinterpretation of data is accomplished is through a `union` which may be used to interpret the same piece of memory in multiple ways. If the use of the union members is not managed carefully, then unexpected and erroneous results may occur.

C allows the use of pointers to memory so that an integer pointer could be used to manipulate character data. This could lead to a mistake in the logic that is used to interpret the data leading to unexpected and erroneous results.

C.48.3 Avoiding the vulnerability or mitigating its effects

- Avoid the use of unions as it is relatively easy for there to exist an unexpected program flow that leads to a misinterpretation of the union data.

C.48.4 Implications for standardization

Future standardization efforts should consider:

- Deprecating unions. The primary reason for the use of unions to save memory has been diminished considerably as memory has become cheaper and more available. Unions are not statically type safe and are historically known to be a common source of errors, leading to many C programming guidelines specifically prohibiting the use of unions.

C.49 Memory Leak [XYL]

C.49.1 Terminology and features

C.49.2 Description of vulnerability

C can allow memory leaks as many programs use dynamically allocated memory. C relies on manual memory management rather than a built in garbage collector primarily since automated memory management can be unpredictable, impact performance and is limited in its ability to detect unused memory such as memory that is still referenced by a pointer, but is never used.

Memory is dynamically allocated in C using the library calls `malloc()`, `calloc()`, and `realloc()`. When the program no longer needs the dynamically allocated memory, it can be released using the library call `free()`. Should there be a flaw in the logic of the program, memory continues to be allocated but is not freed when it is no longer needed. A common situation is where memory is allocated while in a function, the memory is not freed before the exit from the function and the lifetime of the pointer to the memory has ended upon exit from the function.

C.49.3 Avoiding the vulnerability or mitigating its effects

- Use debugging tools such as leak detectors to help identify unreachable memory.
- Allocate and free memory in the same module and at the same level of abstraction to make it easier to determine when and if an allocated block of memory has been freed.
- Use `realloc()` only to resize dynamically allocated arrays.
- Use garbage collectors that are available to replace the usual C library calls for dynamic memory allocation which allocate memory to allow memory to be recycled when it is no longer reachable. The use of garbage collectors may not be acceptable for some applications as the delay introduced when the allocator reclaims memory may be noticeable or even objectionable leading to performance degradation.

C.49.4 Implications for standardization

Future standardization efforts should consider:

None

C.50 Argument Passing to Library Functions [TRJ]

C.50.1 Terminology and features

C.50.2 Description of vulnerability

Parameter passing in C is either pass by reference or pass by value. There isn't a guarantee that the values being passed will be verified by either the calling or receiving functions. So values outside of the assumed range may be received by a function resulting in a potential vulnerability.

A parameter may be received by a function that was assumed to be within a particular range and then an operation or series of operations is performed using the value of the parameter resulting in unanticipated results and even a potential vulnerability.

C.50.3 Avoiding the vulnerability or mitigating its effects

- Do not make assumptions about the values of parameters.
- Do not assume that the calling or receiving function will be range checking a parameter. It is always safest to not make any assumptions about parameters used in C libraries. Because performance is sometimes cited as a reason to use C, parameter checking in both the calling and receiving functions is considered a waste of time. Since the calling routine may have better knowledge of the values a parameter can hold, it may be considered the better place for checks to be made as there are times when a parameter doesn't need to be checked since other factors may limit its possible values. However, since the receiving routine understands how the parameter will be used and it is good practice to check all inputs, it makes sense for the receiving routine to check the value of parameters. Therefore, in C it is difficult to create a blanket statement as to where the parameter checks should be made and as a result, parameter checks are recommended in both the calling and receiving routines unless knowledge about the calling or receiving routines dictates that this isn't needed.

C.50.4 Implications for standardization

Future standardization efforts should consider:

- Creating a recognizable naming standard for routines such that one version of a library does parameter checking to the extent possible and another version does no parameter checking. The first version would be considered safer and more secure and the second could be used in certain situations where performance is key and the checking is assumed to be done in the calling routine. A naming standard could be made such that the library that does parameter checking could be named as usual, say "library_xyz" and an equivalent version that does not do checking could have a "_p" appended, such as "library_xyz_p". Without a naming standard such as this, a considerable number of wasted cycles will be conducted doing a double check of parameters or even worse, no checking will be done in both the calling and receiving routines as each is assuming the other is doing the checking.

C.51 Dynamically-linked Code and Self-modifying Code [NYY]

C.51.1 Terminology and features

C.51.2 Description of vulnerability

Most loaders allow dynamically linked libraries also known as shared libraries. Code is designed and tested using a suite of shared libraries which are loaded at execution time. The process of linking and loading is outside the scope of the C standard.

C can allow self-modifying code. In C there isn't a distinction between data space and code space, executable commands can be altered as desired during the execution of the program. Although self-modifying code may be easy to do in C, it can be difficult to understand, test and fix leading to potential vulnerabilities in the code.

Self-modifying code can be done intentionally in C to obfuscate the effect of a program or in some special situations to increase performance. Because of the ease with which executable code can be modified in C, accidental (or maliciously intentional) modification of C code can occur if pointers are misdirected to modify code space instead of data space or code is executed in data space. Accidental modification usually leads to a program crash. Intentional modification can also lead to a program crash, but used in conjunction with other vulnerabilities can lead to more serious problems that affect the entire host.

C.51.3 Avoiding the vulnerability or mitigating its effects

- Use signatures to verify that the shared libraries used are identical to the libraries with which the code was tested.

- Do not use self-modifying code except in rare instances. In those rare instances, self-modifying code in C can and should be constrained to a particular section of the code and well commented.

C.51.4 Implications for standardization

Future standardization efforts should consider:

- Standardizing on an easy to use signature mechanism for libraries. Standard C libraries should be signed to allow for verification.

C.52 Library Signature [NSQ]

C.52.1 Terminology and features

C.52.2 Description of vulnerability

Integrating C and another language into a single executable relies on knowledge of how to interface the function calls, argument lists and data structures so that symbols match in the object code during linking. Byte alignments can be a source of data corruption.

For instance, when calling Fortran from C, several issues arise. Neither C nor Fortran check for mismatch argument types or even the number of arguments. C passes arguments by value and Fortran passes arguments by reference, so addresses must be passed to Fortran rather than values in the argument list. Multidimensional arrays in C are stored in row major order, whereas Fortran stores them in column major order. Strings in C are terminated by a null character, whereas Fortran uses the declared length of a string. These are just some of the issues that arise when calling Fortran programs from C. Each language has its differences with C, so different issues arise with each interface.

Writing a library wrapper is the traditional way of interfacing with code from another language. However, this can be quite tedious and error prone.

C.52.3 Avoiding the vulnerability or mitigating its effects

- Use a tool, if possible, to automatically create the interface wrappers.
- Minimize the use of those issues known to be error prone when interfacing from C, such as passing character strings, passing multi-dimensional arrays to a column major language, interfacing with other parameter formats such as call by reference or name and receiving return codes.

C.52.4 Implications for standardization

Future standardization efforts should consider:

None

C.53 Unanticipated Exceptions from Library Routines [HJW]

C.53.1 Terminology and features

C.53.2 Description of vulnerability

Calling software routines produced outside of the control of the main application developer puts all of the code at the mercy of the called routines. An unanticipated exception generated from a library routine could have

devastating consequences.

C.53.3 Avoiding the vulnerability or mitigating its effects

- Check the values of parameters to ensure appropriate values are passed to libraries in order to reduce or eliminate the chance of an unanticipated exception

C.53.4 Implications for standardization

Future standardization efforts should consider:

None

DRAFT